

True-concurrency Probabilistic Models: Markov Nets and a Law of Large Numbers

Samy Abbes* Albert Benveniste†

2006

Abstract

We introduce the model of Markov nets, a probabilistic extension of safe Petri nets under the true-concurrency semantics—this means that traces, not firing sequences, are given a probability. This model builds upon our previous work on probabilistic event structures. We use the notion of branching cell for event structures and show that the latter provides the adequate notion of local state, for nets. We prove a Law of Large Numbers (LLN) for Markov nets, which constitutes the main contribution of the paper. This LLN allows characterizing in a quantitative way the asymptotic behavior of Markov nets.

Keywords: Probabilistic event structure, probabilistic Petri net, true-concurrency, probability

1 Introduction

This paper studies the model of *Markov nets*, a probabilistic model of safe Petri nets under true-concurrency semantics. This means that traces (or configurations), not firing sequences, are given a probability.

True-concurrency models and paradigms have been extensively studied since the pioneering works [1, 2]. Such models regard as equivalent executions sequences that differ only via interleaving of concurrent events. They handle directly resulting equivalence classes, also called traces or configurations. True-concurrency semantics differs thus from the more traditional interleaving semantics, where sequences themselves are the subject of study.

While interleaving and true-concurrency theories clearly differ in their mathematics, one could argue that the difference is less obvious from practical standpoint. For example, for systems involving concurrency, interleaving semantics can still be used efficiently, e.g., by picking exactly one representative per equivalence class and discarding the other. This technique has, for instance, been used under the name of “partial order reduction” in the Spin model checking tool¹, see [3]. However, the two

*LIAFA/Université Paris 7 Denis Diderot

†IRISA/INRIA, Rennes, France

¹<http://spinroot.com/>

approaches dramatically differ, both mathematically and practically, when it comes to equipping systems with probabilities. The random outcomes are either the individual maximal sequential executions—for the interleaving semantics—or the maximal traces—for the true-concurrency semantics. The associated space of maximal executions, either sequences or traces, is the relevant space on which a probability measure has to be defined. Markovian properties for instance have then different formulations in the two situations. Hence referring to interleaving sequences as a technical artifact is no longer possible when dealing with true-concurrency probabilistic models.

For situations in which executions should not be distinguished when they differ only via interleaving, true-concurrency probabilistic models are the ones to be considered. Applications such as diagnosis of large distributed systems (e.g., telecommunications networks or services) calls for true-concurrency semantics: alarms or observations are collected throughout the network and sent to the supervisor(s) where they are stored as “logs”. Hence the fine ordering of alarms in logs has only limited meaning and thus logs should rather be confronted to candidate traces of the distributed systems in order to explain the observed behavior. The same holds for the distributed observation of distributed systems in a more general setting. See [4, 5] for detailed explanations.

In this paper, we build on our previous work [6] on probabilistic event structures to randomize Petri nets. Probabilistic event structures are prime event structures whose space of maximal configurations is equipped with a probability. We have defined in [6] a proper notion of “local state” for event structures that proved relevant for their randomization. Such local states are sub-event structures that localize the sources of conflict, we call them *branching cells*. Local probabilities attached to local states were shown to combine together into a *distributed product*, thus defining a global probability from the local ones. The main feature of this randomization is to respect the principle that *concurrent processes should be independent in the probabilistic sense*.

On the other hand, Petri nets, and especially safe Petri nets, are given a true-concurrency semantics in terms of event structures through the canonical operation of *unfolding*. Thus, with our work [6] we were quite ready to randomize Petri nets by using our distributed product. For event structures arising from Petri net unfoldings, local states that are isomorphic when seen as event structures labeled by transitions from the net, are considered equivalent. Corresponding equivalence classes are shown to be only finitely many. Equipping equivalent local states with the same local probabilities defines *Markov nets*.

Since equivalence classes of local states are only finitely many, some of them will be reached infinitely often. How often is quantitatively assessed by the novel *Law of Large Numbers (LLN)* we develop in this paper for Markov nets. The LLN can be seen as a probabilistic reinforcement of the central notion of fairness. It is the fundamental basis for the use of probabilistic models in statistics.

Recall that the classical LLN for a sequence $(X_k)_{k>0}$ of independent random variables with identical distribution μ states as follows: for every nonnegative, real-valued, function f defined on the state space of X_k , the empirical means $\frac{1}{n} \sum_{k=1}^n f(X_k)$ converge with probability 1 to the mean value $\int f(x)\mu(dx)$. This LLN extends to (recurrent) Markov chains, with μ being the invariant measure of the chain. For Markov nets,

the statement of a LLN is by itself doubly challenging, since:

1. there is no global time index k in the true-concurrency framework, and
2. it is unclear what the state space should be.

Regarding the second point, the normal guess that the state should be the configuration of tokens in the places of the net does not work in general—places are not the right notion of local state in the probabilistic context. We will indeed show that our new notion of “local state” is adequate for a statement of the LLN. Regarding the first point, we will see that the “time elapsed while scanning a configuration” is properly measured by the number of local states traversed.

A third challenge is that, since configurations exhibit concurrency, the “progress of time” while scanning a configuration is not clearly defined: the different processes composing the configuration may progress freely as long as they do not need to synchronize. Our LLN requires that empirical means converge whatever way these different processes progress, and that the resulting limit should be unique.

Not every Markov net can obey such a LLN: in a net composed of two non interacting subnets, the two subnets can progress freely and no single “life time” can bound this progress. For such nets lacking synchronization, our LLN cannot hold. We thus finely characterize how much lack of synchronization can be tolerated, while still having the LLN valid: this is expressed as an integrability condition on a certain random variable that measures the lack of synchronization of the system. This condition is trivially satisfied for Markov chains, since Markov chains do not present any concurrency and thus “maximally synchronize”. In turn, we expect—although it is not covered by this paper—that the condition is not really a restriction, since a net should be decomposed into subnets, each of them satisfying our synchronization condition. That would be a synchronization decomposition, analogous to the usual decomposition of Markov chains through connected components.

To keep this paper self-contained, the key results of [6] concerning local states of event structures and distributed probabilities are recalled. The paper is organized as follows: In §2, we introduce on a few toy examples the kind of randomization for Petri nets we consider in this paper. The LLN for true-concurrency systems is also informally discussed. This section is intended to both 1) researchers used to the interleaving semantics for probabilistic systems and who wish to have a short introduction to its true-concurrency counterpart, and 2) researchers comfortable with the true-concurrency semantics but not with its probabilistic counterpart. Then, we collect in §3 some basic notions concerning event structures and net unfoldings, and fix the notations. The non-standard notions of stopping prefixes and branching cells (the notion formalizing our “local states”) are borrowed from [6] and collected in §4. In §5, we recall the basic notion concerning probabilistic event structures as well as the construction of distributed probabilities. Markov nets are introduced in §6, and we state the LLN in §7. The proof of the LLN is detailed in §8. Related work is discussed in §9. In Appendix A, we recall the classical statements of the LLN that we use.

2 Illustrative Examples

This section presents some toy examples to illustrate the issues encountered when randomizing traces for a Petri net, and the solutions we propose. We first recall basic concepts of trace theory. We explain on simple examples why *dynamic clusters* of nets must be considered, and their role in the randomization of nets—a more detailed discussion on this topic is found in § 5.4.. We also informally describe the Law of Large Numbers in this context.

True Concurrency and Traces. According to the true-concurrency approach to safe Petri nets, we do not distinguish a firing sequence (\dots, t, t', \dots) from the firing sequence obtained by exchanging transitions t and t' , whenever t and t' are transitions of the net that share no common resource. The equivalence classes—after taking the reflexive transitive closure of the above relation—are called the traces of the net [2, 1]. Figure 1 depicts the example of a safe Petri net. In this example, the firing sequence (bce) is equivalent to the sequence (cbe) since transitions b and c do not share common resources.




Figure 1: *Illustrating the true-concurrency semantics.*

Firing sequences of a safe Petri net are ordered by the prefix relation on words. This relation induces in turn an ordering on traces, so that trace σ precedes σ' if and only if there are firing sequences s and s' , with s a prefix of s' , representing σ and σ' , respectively. The partial order on traces for the net of Figure 1 is depicted in Figure 2. On the other hand, each trace is seen itself as a partially ordered multi-set (*pomset*) of transitions, where two transitions are not comparable when they can be exchanged according to the above rule. In the language of event structures, the elements that compose the pomset are called *events*. This is illustrated in Figure 2, (2), for the trace $(bce) = (cbe)$.

Randomizing Maximal Traces. Traces are partially ordered sets of events. Therefore, classical approaches from the area of stochastic processes do not apply, since they assume a totally ordered time. We thus first discuss how traces should be randomized.




Figure 2: (1) Ordering of the traces of the net depicted in Figure 1. (2) The trace $(bce) = (cbe)$ as a partial order of labeled events.

We randomize a safe Petri net by defining a probability on its set of maximal traces. That is, we implicitly agree that *no blocking other than intrinsic deadlocks of the net shall appear due to randomization*. Allowing the probability to weight non-maximal traces could be relevant for modeling reasons, for example, to express the possibility of failure due to a program crash. In the classical study of finite Markov chains, we indeed find this feature under the name of “killed processes” [7, p.25]. It is also known that killed trajectories can be made maximal by adding an additional dummy state in the state space, called the “cemetery”. The same can be performed for nets, so we only focus on maximal traces throughout this paper. For the example depicted in Figure 1, maximal traces are the extremal bottom points of Figure 2 (1), i.e., (bd) , (ad) , (ac) and (bce) . Making the net probabilistic amounts thus to defining a probability \mathbb{P} such that:

$$\mathbb{P}(bd) + \mathbb{P}(ad) + \mathbb{P}(ac) + \mathbb{P}(bce) = 1.$$

Observe that, if we sum these probabilities, not on traces but on firings sequences, the result will exceed 1. This means that considering true-concurrency semantics has a significant impact on the construction of probabilities.

Concurrency and Probability. Our constant philosophy in the construction of probabilistic nets is the following: as much as possible, parallel processes shall be made independent in the probabilistic sense. When considering a Petri net as a distributed system with a distributed state, this requirement is quite natural. In turn, synchronization is a major source of difficulty. Indeed, when components do not interact at all, we can simply randomize each component separately, and make them probabilistically independent by decree. In general, however, processes are “parallel” for a limited amount of “time”, then synchronize, which results in breaking the parallelism. We shall give a precise formulation of the latter claim, by defining a decomposition of processes through locally parallel components. With this decomposition of processes at hand, we shall proceed with their randomization, by enforcing probabilistic independence of locally parallel components. This is not a trivial task, but is successfully achieved by

using classical tools from Measure theory. The resulting theory is presented in detail in [6], for the model of event structures.

Dynamic Clusters. Consider first the two nets depicted in Figure 3, (a) and (b). Net in (a) offers three possibilities: either firing transition a , or firing transition b , or firing no transition at all. Since we only consider maximal traces, as explained above, we deny to this net the right of doing nothing, so that it must eventually fire transition a or b . This is a simple coin tossing, that occurs with a certain probability, say $\mu_1(a)$ versus $\mu_1(b)$.




Figure 3: *Illustrating local parallelism and synchronization of safe Petri nets.*

Consider next the net of Figure 3, (b). Again, since we consider only maximal traces for randomization, the net has exactly two possibilities: either firing transition d , or firing concurrently transitions c and e . We have again a coin tossing. This time, however, the tossing does not involve single transitions, but rather certain groups of transitions. We shall thus weight the occurrence of (ce) with some probability $\mu_2(ce)$, while $\mu_2(d)$ is the probability that (d) occurs instead of (ce) . Remark that, although transitions c and e are concurrent, they are not independent, since c occurs if and only if e occurs.

Let us now investigate the effect of synchronization. The net of Figure 3 (c) collects the two previous nets and adds some new elements whose execution depends on the previous execution of the first two nets. In this discussion, we shall refer to the first two nets, considered separately, as *local nets*, and to their executions as *local executions*. In contrast, the net of Figure 3 (c) is referred to as the *global net*. Because the two local nets do not interact at all, the local execution of one of them does not disturb the execution of the other one. Hence the local executions are *concurrent* and do not interact.

Suppose that each local net is driven by some random local agent, and that these local agents do not communicate with each other. It is then natural, from the probabilistic point of view, to assume that the actions of local agents are independent in the probabilistic sense. Observe that a maximal trace of the global net induces, by restriction, a maximal trace of each local net. This is important, since we have only specified how to randomize the maximal traces of the two local nets.

Consequently, we know at this point how to randomize the “beginning” of the executions of the global net. If μ_1 and μ_2 denote respectively the probabilities attached to each local net, so that

$$\mu_1(a) + \mu_1(b) = 1, \quad \mu_2(ce) + \mu_2(d) = 1,$$

and if \mathbb{P} denotes the global probability constructed so far, we have for example:

$$\mathbb{P}(ace) = \mu_1(a) \times \mu_2(ce), \quad \mathbb{P}(bce) = \mu_1(b) \times \mu_2(ce).$$

In the above equation, the left members actually concern maximal executions of the global net: they must be understood as, for example: “ $\mathbb{P}(ace)$ is the probability that a maximal execution contains (ace) ”.

What happens next? Consider for example the case of (ace) having fired. The resulting net is depicted in Figure 4 (a). The transitions that are not enabled are depicted with dashes, and we omit the transitions that have already fired. The only enabled transitions are transitions h and i . We shall thus consider some probability μ_3 , that describes the weight of h against i , *in the context of h and i competing alone*. By the chain rule, the complete probabilities of $(aceh)$ and of $(acei)$ can now be computed by:

$$\mathbb{P}(aceh) = \mu_1(a) \times \mu_2(ce) \times \mu_3(h), \quad \mathbb{P}(acei) = \mu_1(a) \times \mu_2(ce) \times \mu_3(i).$$

Instead of (a) and (ce) being the executions of the two first local nets, assume that, for example, (b) and (ce) have fired. The resulting net is depicted in Figure 4 (b). Now all transitions f , g , h and i are enabled. Hence h and i still compete for firing, but they do not compete in the same context as previously observed, since f and g also take part to the competition. We shall thus consider yet another local probability μ_4 , that randomizes the maximal traces of the net depicted in Figure 4 (b), i.e., such that:

$$\mu_4(fh) + \mu_4(fi) + \mu_4(gi) = 1.$$

We again use the chain rule to compute the probability of, say, $(bcefh)$:

$$\mathbb{P}(bcefh) = \mu_1(b) \times \mu_2(ce) \times \mu_4(fh).$$




Figure 4: Nets resulting from different partial executions of the net depicted in Figure 3 (c).

Finally, in case of (ad), respectively (bd), firing in the first two local nets, we would consider also the additional local net consisting of transition i only, respectively the additional local nets consisting of transition f only and of transition i only. Both nets are necessarily equipped with trivial probabilities, i.e., with probabilities giving weight 1 to the unique possible transition.

With the recursive decomposition of traces described so far, we have reached all maximal configurations of the global net and we know how to compute the probability of each maximal configuration. We leave to the reader as an exercise to check on this example that the probability defined by this way sums indeed to 1 on the set of all maximal configurations of the global net, by using the fact that μ_1 and μ_2 both sum to 1.

The different local nets encountered in the course of all possible executions of the net are called *dynamic clusters*. The global net we have studied has 6 dynamic clusters: the two nets of Figure 3 (a) and (b), the subnet obtained from Figure 4 (a) by keeping only transitions h and i , the one obtained from Figure 4 (b) by keeping only transitions f , g , h and i , and finally the two trivial nets with single transitions f and i .

The decomposition through these subnets is indeed *dynamic*, since a same transition may occur in different dynamic clusters, according to the context. This was the case, for example, for the transitions h and i . This fact conforms with the intuition that concurrent systems shall not be statically decomposed, but may split and join in different manners, according to the actual execution. We summarize what we have obtained so far:

1. We randomize maximal traces of safe Petri nets.
2. Maximal traces are decomposed as the juxtaposition of maximal traces of dynamic clusters. Hence dynamic clusters appear as *local states*, on which local actions operate by means of local maximal traces. These are randomized, according to some fixed probability law.
3. *Concurrent* dynamic clusters do not interfere with each other. Therefore, concurrent dynamic clusters are made independent in the probabilistic sense, which

eventually brings a global randomization of the system built up from the local randomizations.

The probabilistic systems constructed in this way are called *Markov nets*.

Regarding the Law of Large Numbers. Recall the classical LLN for a sequence $(X_k)_{k>0}$ of independent random variables with identical distribution μ : for every non-negative, real-valued, function f defined on the state space of X_k , the empirical means $\frac{1}{n} \sum_{k=1}^n f(X_k)$ converge with probability 1 to the mean value $\int f(x)\mu(dx)$. This LLN extends to (recurrent) Markov chains, with μ being the invariant measure of the chain. We shall now discuss the LLN for Markov nets. Of course, such a LLN is relevant only for nets having infinite configurations.

An example of such a net is shown in Figure 5 and its dynamic clusters are depicted in Figure 6. Note that these clusters overlap, revealing their dynamic nature. For example, transition d belongs to both clusters s_4 and s_5 ; in fact, when a trace v traverses d , then it traverses either cluster s_4 or cluster s_5 but it cannot traverse both.

As observed in the introduction, the first difficulty consists in finding the adequate notion of state, for Markov nets. Since dynamic clusters have proved to be adequate so far as candidates for local states, why wouldn't we try to push a little bit more the analogy? Therefore, our "state functions" will be real-valued functions defined on the set Σ of dynamical clusters. The next issue is that of finding the counterpart of the time index k in empirical means $\frac{1}{n} \sum_{k=1}^n f(X_k)$. We regard k as indexing the successive outcomes of the random choices made while drawing the considered trajectory. Therefore, a natural counterpart of k is the number of dynamic clusters traversed while constructing the considered trace.




Figure 5: A safe net with infinite executions. Places having the same name are to be superimposed.

The above analysis suggests the following form for the empirical means in our case:

$$M(f, v) = \frac{\sum_{s \in \Sigma : v \text{ traverses } s} f(s)}{\text{number of dynamic clusters } s \text{ traversed by } v}, \quad (1)$$




Figure 6: Dynamic clusters s_1, \dots, s_5 , of the net depicted in Figure 5. The figure depicts the unfolded form of the clusters. Here letters indicate the labels, not the names of the node, and places with same labels are not to be superimposed.

where f is a state function and v is some finite trace. It must be noted that a same cluster is traversed several times by the finite trace v , such traversals are therefore counted with their multiplicities in the above summation \sum_s .

Pick some particular dynamic cluster s_* , and consider (1) with $f(s) = \mathbf{I}_{s_*}(s)$, where $\mathbf{I}_{s_*}(s) = 1$ if $s = s_*$, and $= 0$ otherwise. Then,

$$M(\mathbf{I}_{s_*}, v) = \frac{\text{number of occurrences of } s_* \text{ seen by } v}{\text{number of dynamic clusters } s \text{ traversed by } v}.$$

Assume for a while that, for each maximal trace ω and each dynamic cluster s_* ,

$$\alpha(s_*) = \underset{[\text{def}]}{\lim_{v \nearrow \omega}} M(\mathbf{I}_{s_*}, v) \tag{2}$$

exists and does not depend on ω . Since, for every v , $\sum_s M(\mathbf{I}_s, v) = 1$, the $\alpha(s)$ sum up to 1 and therefore define a probability on the finite set Σ , we call it the *stationary measure* of the Markov net. Then, by linearity, for every nonnegative state function f ,

$$\lim_{v \nearrow \omega} M(f, v) = \sum_{s \in \Sigma} f(s) \alpha(s)$$

exists and does not depend on ω . This provides us with the desired LLN. Thus it is enough proving the LLN for the special case (2).

A direct application of this LLN yields in particular the asymptotic ratio:

$$\lim_{v \nearrow \omega} \frac{\text{number of occurrences of } s \text{ in } v}{\text{number of occurrences of } s' \text{ in } v} = \frac{\alpha(s)}{\alpha(s')},$$

for s, s' two dynamic clusters (see Figure 6 for an illustration of the clusters). Now, a more natural question would be to evaluate the asymptotic ratio:

$$\lim_{v \nearrow \omega} \frac{\text{number of occurrences of } t \text{ in } v}{\text{number of occurrences of } t' \text{ in } v}, \quad (3)$$

for t, t' two transitions of the net. Unfortunately, transitions are not in bijection with dynamic clusters (as shown by Figure 6), and therefore the above asymptotic ratio is not within the scope of our above LLN.

Suppose now we were able to prove a LLN for “extended” state functions of the form $f(s, \omega_s)$, where s is a dynamic cluster as before, and ω_s is a maximal trace of s , seen as an event structure. Corresponding empirical means would be:

$$M(f, v) = \frac{\sum_{s \in \Sigma: v \text{ traverses } s} f(s, \omega_s)}{\text{number of dynamic clusters } s \text{ traversed by } v},$$

where ω_s is the maximal trace of s seen by v when traversing s , and the traversals of s by v are counted with their multiplicities, as in (1). Then,

$$\mathbf{I}_t(s, \omega_s) \stackrel{[\text{def}]}{=} \begin{cases} 1, & \text{if } \omega_s \text{ visits } t \\ 0, & \text{otherwise} \end{cases}$$

would yield an extended state function such that

$$\frac{M(\mathbf{I}_t, v)}{M(\mathbf{I}_{t'}, v)} = \frac{\text{number of occurrences of } t \text{ in } v}{\text{number of occurrences of } t' \text{ in } v},$$

so that our extended LLN would encompass asymptotic ratios of the form (3). We shall indeed prove such an extended LLN and therefore solve the problem of evaluating asymptotic ratios of the form (3).

So far we have used the expression

$$\lim_{v \nearrow \omega} M(f, v) \quad (4)$$

without care. Indeed, trace v can grow to ω in many different ways. If, for example, we erase, in the net of Figure 5, the places B and E , and the arcs $D \rightarrow e$ and $e \rightarrow A$, then the two remaining non interacting subnets can progress freely. And it is unclear whether the limit (4) is well defined, since it could very well depend on the way the two components of v grow. Not surprisingly, the LLN we shall establish assumes that the different local processes of the net do synchronize “frequently enough”. In turn, we will show that, with this assumption, the limit (4) is well defined and does not depend on the way v can grow to ω .

3 Prerequisites on Event Structures and Unfoldings

In this section we recall basic definitions concerning event structures, Petri nets, and their unfoldings.

3.1 Event Structures

Throughout this paper, the term “event structure” means *prime* event structure.

An event structure [1] is a triple $(E, \preceq, \#)$ satisfying the following properties. (E, \preceq) is a partial order, at most countable, whose elements are called events. $\#$ is the *conflict* relation on E ; it is a binary, symmetric and irreflexive relation, that satisfies the following axiom: $\forall x, y, z \in E, x\#y$ and $y \preceq z$ together imply $x\#z$. We also assume that, for every $e \in E$, the downward closure of e , defined by $\lceil e \rceil = \{f \in E : f \preceq e\}$ is a finite subset of E . We identify, with a slight abuse of notations, the set E and the event structure $(E, \preceq, \#)$. Finally, the *concurrency relation* is the binary relation on E , denoted by co , and defined by $\text{co} = (E \times E) \setminus (\# \cup \preceq \cup \succeq)$. We say that E is a *tree of events*, or shortly, a *tree*, if $\text{co} = \emptyset$.

A subset A of E is said to be a *prefix* if it is downward closed, i.e., if $\lceil e \rceil \subseteq A$ for all $e \in A$. A prefix v is called a *configuration* of E if it is *conflict-free*, i.e., if $\# \cap (v \times v) = \emptyset$. Configurations are partially ordered by inclusion, and we denote by \mathcal{V}_E the poset of *finite* configurations of E . Two configurations are said to be *compatible* if their set-theoretic union is conflict-free. We denote by Ω_E the set of maximal configurations of E —this set is nonempty according to Zorn lemma, since any chain of configurations has an upper bound. If no confusion can occur, we simply use the notation Ω ; this notation is indeed reminiscent from probability theory, the reason will be explained in §5.

A subset $F \subseteq E$ defines a sub-event structure $(F, \preceq_F, \#_F)$ of E with causality and conflict relations inherited by:

$$\preceq_F = \preceq \cap (F \times F), \quad \#_F = \# \cap (F \times F),$$

and we shall freely write F, \mathcal{V}_F , and Ω_F to denote this event structure and its set of finite and maximal configurations, respectively.

3.2 Petri nets

Petri nets represent finite systems with concurrency properties. A Petri net is a tuple $\mathcal{N} = (P, T, F, m_0)$, where P and T are two disjoint finite sets, respectively of *places* and of *transitions*, F is the *flow relation*, defined as a subset of $P \times T \cup T \times P$. Observe that the triple (P, T, F) defines a graph, where arcs from places point towards transitions, and arcs from transitions point towards places. If t is a transition, the *preset* of t is defined as the set of places p such that $(p, t) \in F$, and is denoted $\bullet t$. Similarly, the *postset* of t is the set of places p such that $(t, p) \in F$. A *marking* of \mathcal{N} is defined as an integer-valued mapping $m : P \rightarrow \mathbb{N}$; symbol m_0 above denotes the *initial marking* of the net, meaning that each place p is initially filled with $m(p)$ tokens.

The *Petri net rule* describes the dynamics of the net. Let m be a marking of \mathcal{N} , and let t be a transition of \mathcal{N} . We say that t is *enabled* by marking m if $m(p) \geq 1$ for all $p \in \bullet t$. In this case, the *firing* of t transforms m into the marking m' defined by $m'(p) = m(p)$ if $p \notin \bullet t \cup t^\bullet$, and by

$$m'(p) = m(p) - 1 \quad \text{if } p \in \bullet t, \quad m'(p) = m(p) + 1 \quad \text{if } p \in t^\bullet.$$

We note this transformation rule by $m \xrightarrow{t} m'$. We say that a sequence of transitions (t_1, \dots, t_n) is enabled from m whenever, for some markings m_1, \dots, m_n we have

$$m \xrightarrow{t_1} m_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} m_n.$$

In this case, marking m_n is said to be *reachable from m* . We simply talk about *reachable markings* if $m = m_0$, and we say that (t_1, \dots, t_n) is a *firing sequence* of the net. We also consider the case of empty sequences of transitions, which have no effect on markings.

In the rest of the paper, we will only consider the restricted class of *safe or 1-bounded Petri nets*, defined as follows: a Petri net is said to be *safe* if any reachable marking m satisfies: $\forall p \in P, m(p) \leq 1$. We then identify m with a subset of P , the characteristic function of which is m , and in particular the initial marking m_0 is identified with an initial subset of places.

3.3 Traces and Unfoldings

A particular feature of Petri nets is their concurrency properties, highlighted by the so-called true-concurrency semantics we explain now. Consider the binary relation \sim defined on firing sequences of a 1-bounded Petri net \mathcal{N} by

$$(\dots, s, t, \dots) \sim (\dots, t, s, \dots) \quad \text{whenever } \bullet t \cap \bullet s = \emptyset.$$

Hence, two consecutive transitions in the sequence can be exchanged when they do not share places in their presets. Also consider the reflexive transitive closure of \sim , that we still denote by \sim . This defines an equivalence relation on firing sequences, the equivalence classes of which are called *traces*. Traces inherit a partial order from the prefix relation on firing sequences. If u and v are two equivalent firing sequences from some marking m , then u and v reach the same marking. Hence the marking reached by some trace is well defined.

The unfolding theory [1] states that traces can be represented as the finite configurations of some labeled event structure, canonically attached to the net. Define first a *X-labeled event structure*, for X a set, as a pair (E, λ) , where E is an event structure and $\lambda : E \rightarrow X$ is a mapping, called labeling. The *unfolding* of $\mathcal{N} = (P, T, F, m_0)$ is a *T-labeled event structure* (E, λ) characterized as follows: Say that a sequence (e_1, \dots, e_n) is a *firing sequence* of E if e_1, \dots, e_n are pairwise distinct events of E , such that $\{e_1, \dots, e_j\}$ is a configuration of E for all $j = 1, \dots, n$. Firing sequences of E are partially ordered by the prefix relation. Then the mapping λ , extended in the obvious way to firing sequences of E , induces an isomorphism of partial orders between firing sequences of E and firing sequences of \mathcal{N} . Moreover, if $s = (e_1, \dots, e_n)$ and $s' = (e'_1, \dots, e'_{n'})$ are two firing sequences of E , then $\lambda(s) \sim \lambda(s')$ if and only if $n = n'$ and the equality of sets $\{e_1, \dots, e_n\} = \{e'_1, \dots, e'_{n'}\}$ holds. The existence of unfoldings for safe Petri nets is shown in [1], together with their uniqueness up to unique isomorphisms.

This shows that traces are adequately represented by finite configurations of the unfolding. In the remaining of the paper, this is the only representation of traces we will consider.

3.4 Future of a Configuration

For v a finite or infinite configuration of event structure E , we consider the following subset of E :

$$E^v =_{[\text{def}]} \{e \in E \setminus v : \forall e' \in v, \neg(e\#e)\}$$

We call the associated sub-event structure the *future* of v . It is clear that the map:

$$w \mapsto v \cup w \tag{5}$$

is one-to-one and onto, from the set of configurations of E^v to the set of configurations of E that contains v . The map (5) describes the *concatenation* of v with configurations of E^v . This map is also order-preserving, so that it maps maximal elements to maximal elements; that is, the set of maximal configurations of E^v is mapped onto the set of maximal configurations of E that contain v .

In order to distinguish $v \cup w$, for $v \in \mathcal{V}_E$ and $w \in \mathcal{V}_{E^v}$ from the usual set-theoretic union of compatible configurations of E , we introduce the following special notations for the concatenation and for its inverse:

$$\begin{aligned} v \oplus w &=_{[\text{def}]} v \cup w, & \text{defined for } v \in \mathcal{V}_E \text{ and } w \in \mathcal{V}_{E^v}, \\ u \ominus v &=_{[\text{def}]} u \setminus v, & \text{defined for } v \in \mathcal{V}_E \text{ and } u \in \mathcal{V}_E \text{ such that } v \subseteq u, \end{aligned} \tag{6}$$

so that we have, whenever these are well defined:

$$v \oplus w \in \mathcal{V}_E, \quad u \ominus v \in \mathcal{V}_{E^v}.$$

In the context of net unfoldings, the operation \oplus defined above corresponds to the concatenation of traces of the net (inherited from the concatenation of firing sequences). Moreover, if u is a finite configuration of the unfolding (E, λ) of a net $\mathcal{N} = (P, T, F, m_0)$, and if m is the marking reached by the trace $\lambda(u)$ associated with u , then (E^u, λ) is the unfolding of the net $\mathcal{N}' = (P, T, F, m)$ obtained from \mathcal{N} by only changing the initial marking.

Most material introduced in this paper regarding event structures is dedicated to event structures obtained by unfolding a safe Petri net (see below §3.3). Although we will first consider abstract event structures and randomize them as such, it is legitimate to assume that any event structure E we consider satisfies the following assumption, which means that any configuration enables only finitely many events:

Assumption 1 *For every $v \in \mathcal{V}_E$, the set $\text{Min}(E^v)$, collecting all minimal events of E^v , is bounded by some constant.*

4 Branching Cells

This section introduces branching cells, a notion introduced in [8, 6] to formalize the notion of “local state”. Our motivations for introducing this notion are mainly probabilistic. Still, the notion of local state for event structures and nets that branching cells provide is of interest *per se*. Branching cells and their properties are introduced *a priori* and we postpone to §5.4 the discussion on why we chose *this* particular definition.

4.1 Stopping Prefixes

A central concept in defining probabilities is the notion of choice. Choice is therefore a key concept in this paper. It is captured for event structures by the notion of minimal conflict we recall next. The *minimal conflict* relation $\#_m$ in an event structure E is defined by:

$$\forall e, e' \in E, \quad e \#_m e' \iff ([e] \times [e']) \cap \# = \{(e, e'), (e', e)\}.$$

In case of the unfolding of a 1-bounded Petri net, two transitions e, e' with $e \#_m e'$ can be *simultaneously enabled*, and are labeled by two transitions t and t' sharing common places in their presets.

Definition 4.1 (stopping prefix) *A subset $B \subseteq E$ is called a stopping prefix of E if:*

1. B is a prefix of E ;
2. B is closed under minimal conflict.

Stopping prefixes form a complete lattice, with \emptyset and E as minimal and maximal elements. Probabilistic constructions consist in randomizing choices and should be therefore based on stopping prefixes. Hence the following notion is natural in this context:

Definition 4.2 (locally finite event structure) *E is called locally finite if for each event e of E , there exists a finite stopping prefix containing e .*

Locally finite event structures have not been considered by authors so far. The following condition is implicitly assumed throughout this paper:

Assumption 2 *E is locally finite.*

It is easily checked that, if v is any configuration of E , e an event of E^v , and B a finite stopping prefix containing e , then $B \cap E^v$ is a finite stopping prefix of E^v containing e . As a consequence, every future E^v is locally finite. Stopping prefixes satisfy the following property (see [9, Ch.3, I-3.1]):

Lemma 4.3 *If B is a stopping prefix of E , then:*

$$\Omega_B = \{\omega \cap B \mid \omega \in \Omega_E\}. \quad (7)$$

This property, which is not satisfied by arbitrary prefixes (draw an example!), motivates the following definition:

Definition 4.4 (stopped configurations) *A configuration v is called a stopped configuration of E if there is a stopping prefix B such that $v \in \Omega_B$.*

REMARK. Use Lemma 4.3 and the fact that E is locally finite to show that a stopped configuration v is finite if and only if there is a *finite* stopping prefix B such that $v \in \Omega_B$.

4.2 Branching Cells and R -stopped Configurations

We seek for the following two-steps procedure for constructing probabilistic event structures: 1) consider stopped configurations as “elementary process” for randomization, and 2) use the concatenation of stopped configurations in order to randomize processes in an incremental way.

Unfortunately, the class of stopped configurations is not closed under concatenation. That is, if v is a stopped configuration of E , and w a stopped configuration of the future E^v , then $v \oplus w$ is not stopped in E in general² (see an example at the end of §4.3). This is why we consider the closure of stopped configurations under concatenation. The configurations reached by this way are called *recursively stopped* (R -stopped for short).

Definition 4.5 (R -stopped configurations) We denote by $\overline{\mathcal{W}}_E$ the smallest class of configurations with the following properties:

1. $\emptyset \in \overline{\mathcal{W}}_E$;
2. $\overline{\mathcal{W}}_E$ is closed under concatenation with finite stopped configurations; that is to say, if $u \in \overline{\mathcal{W}}_E$ and if v is a finite stopped configuration of E^u , then $u \oplus v \in \overline{\mathcal{W}}_E$;
3. $\overline{\mathcal{W}}_E$ is closed under supremum of nondecreasing sequences.

Configurations in $\overline{\mathcal{W}}_E$ are called R -stopped (for recursively stopped). We use the notation \mathcal{W}_E to denote finite R -stopped configurations of E , or the short notation \mathcal{W} if no confusion can occur.

As a consequence, a configuration v is R -stopped if there exists a (finite or infinite) nondecreasing sequence $(v_n)_{0 \leq n < N}$ of configurations, $N \leq +\infty$, satisfying the following conditions:

1. $v_0 = \emptyset$, $v = \bigcup_{0 \leq n < N} v_n$, and
2. $v_n \ominus v_{n-1}$ is a finite stopped configuration of the future $E^{v_{n-1}}$, for every $0 < n < N$.

Although R -stopped configurations are characterized by the existence of a recursive decomposition through stopped configurations, such a decomposition is by no way unique. We shall however focus on *minimal* decompositions. Branching cells are introduced to this end.

Definition 4.6 (initial stopping prefix, branching cell) A stopping prefix B is called initial in E if \emptyset is the only stopping prefix of E strictly contained in B .

Call branching cell of E any initial stopping prefix of some E^v , for v ranging over \mathcal{W}_E . For $v \in \mathcal{W}$, initial stopping prefixes of E^v are called branching cells enabled at v . The set of branching cells enabled at v is denoted

$$\delta_E(v),$$

²This is particular to systems with concurrency, and more precisely to systems with confusion: for confusion-free event structures, stopped configurations are closed under concatenation.

or simply by $\delta(v)$ when no confusion can occur.

The set of all branching cells of E is denoted by X_E (or simply X) and branching cells are generically denoted by the symbol x .

Notice that, as a particular case, $\delta(\emptyset)$ denotes the set of initial stopping prefixes of E . These are the “first” encountered branching cells. Other branching cells derive from the same construction, but shifted inside the event structure, since we have:

$$\forall v \in \mathcal{W}_E, \quad \delta_E(v) = \delta_{E^v}(\emptyset).$$

Branching cells represent the minimal steps that can be crossed when trying to decompose some given R -stopped configuration. The different branching cells encountered do not depend on the particular decomposition chosen. They constitute thus an *invariant*, characteristic of the given R -stopped configuration. This is the meaning of the following result, that formally collects the different key properties of branching cells (proofs are found in [9, Ch.3], or in [6]).

Theorem 4.7

1. **Existence and finiteness of branching cells.** Every branching cell of E is a finite sub-event structure of E . For every $v \in \mathcal{W}_E$, $\delta_E(v)$ is empty if and only if v is maximal in E .
2. **Concurrent branching cells.** Let v be any finite R -stopped configuration. Different branching cells in $\delta_E(v)$ are disjoint and concurrent, the latter meaning that:

$$\forall x, y \in \delta_E(v), \quad x \neq y \implies \forall (e, f) \in x \times y, \quad e \text{ co } f.$$

This has the two following consequences:

- (a) For $v \in \mathcal{W}_E$, let B be a stopping prefix of E^v of the form $B = \bigcup_{x \in \xi} x$, where ξ is any subset of $\delta_E(v)$. Then configurations and maximal configurations of B respectively decompose as the following products:

$$\mathcal{V}_B = \prod_{x \in \xi} \mathcal{V}_x, \quad \Omega_B = \prod_{x \in \xi} \Omega_x. \quad (8)$$

- (b) For any $v \in \mathcal{W}_E$, $\delta_E(v)$ is finite (i.e., there are finitely many concurrent branching cells), and even bounded by some constant K .

3. **Covering map.** For v any R -stopped configuration, there exists a non-decreasing sequence of configurations $(v_n)_{0 \leq n < N}$, with $N \leq +\infty$, and a sequence of branching cells $(x_n)_{0 < n < N}$ with $x_n \in \delta_E(v_n)$ for all n , such that:

- (a) $v_0 = \emptyset$, $v = \bigcup_{0 \leq n < N} v_n$, and
- (b) $v_n \ominus v_{n-1}$ is a maximal configuration of x_n for all $0 < n < N$.

Branching cells $\{x_n, 0 < n < N\}$ are pairwise disjoint. If $(v'_n)_{0 \leq n < N'}$ is another such decomposition, with associated branching cells $(x'_n)_{0 < n < N'}$, then we have the equality of sets:

$$\{x_n, 0 < n < N\} = \{x'_n, 0 < n < N'\}.$$

In particular, $N = N'$. We call covering map of E the map Δ_E defined by:

$$\Delta_E(v) = \{x_n, 0 < n < N\}, \quad (9)$$

which only depends on v and on E .

4. **Covering maps in stopping prefixes.** For any stopping prefix B of E , and for any configuration u of B , u is R -stopped in B if and only if u is R -stopped in E . In this case the covering maps Δ_E and Δ_B satisfy:

$$\Delta_B(v) = \Delta_E(v).$$

In particular, $X_B \subseteq X_E$, with X denoting the sets of branching cells.

5. **Concatenation and subtraction of R -stopped configurations; covering the future.** The class of R -stopped configurations is stable under concatenation and under subtraction:

$$\begin{aligned} \forall u \in \mathcal{W}_E, \quad \forall v \in \mathcal{W}_{E^u}, \quad u \oplus v \in \mathcal{W}_E, \\ \forall u \in \mathcal{W}_E, \quad \forall w \in \mathcal{W}_E, \quad u \subseteq w \Rightarrow w \ominus u \in \mathcal{W}_{E^u}. \end{aligned}$$

For $u \in \mathcal{W}_E$ and $v \in \mathcal{W}_{E^u}$, the covering maps Δ_E and Δ_{E^u} in the future E^u satisfy:

$$\Delta_E(u \oplus v) = \Delta_E(u) \cup \Delta_{E^u}(v), \quad \Delta_E(u) \cap \Delta_{E^u}(v) = \emptyset. \quad (10)$$

In particular, $X_{E^u} \subseteq X_E$ for any $u \in \mathcal{W}_E$. Finally, compatible R -stopped configurations form a lattice.

It must be noted that, except for the property that branching cells are finite, all results stated above remain valid without the local finiteness assumption. Further results requiring local finiteness will be stated in §4.4. First, we detail some examples.

4.3 Examples of Decompositions

For all examples of this paper, we write (abc) to denote the configuration $\{a, b, c\}$. To depict event structures, we use arrows for representing the causality and zigzag arcs for the minimal conflicts, as in Figure 7.

Figure 7 depicts with dashed frames the nonempty stopping prefixes of an event structure E . x and y are the two initial stopping prefixes of E . In this example, configuration (a) is not R -stopped. Indeed, otherwise (a) would be maximal in x , which it is not ((ac) is maximal in x). Hence, *there exists in general finite configurations that are not R -stopped. A fortiori, such configurations are not stopped.*




Figure 7: **Illustrating stopping prefixes and initial stopping prefixes.** Causality is depicted by arrows, minimal conflicts are depicted with zigzag arcs. All nonempty stopping prefixes are shown with dashed frames. x and y are the two initial stopping prefixes of the event structure. Observe that e is a minimal event although it does not belong to any initial stopping prefix.




Figure 8: **Illustrating the decomposition of R -stopped configurations.** (1), an event structure with three maximal configuration $\omega_1 = (ad)$, $\omega_2 = (bce)$ and $\omega_3 = (ace)$. (2), the future of (b) , and (3), the future of (a) . Initial stopping prefixes are depicted with dashed frames.

Still in the event structure of Figure 7, let v be the configuration $v = (aca')$. v is maximal in $B_0 = x \cup y$, hence v is stopped and in particular v is R -stopped. We show that the covering $\Delta_E(v)$ is given by $\Delta_E(v) = \{x, y\}$. Since v is stopped in B_0 , it follows from Lemma 4.3 that $v_x = v \cap x$ is stopped in x . Hence, by point 5 of Theorem 4.7, we have:

$$\Delta_E(v) = \Delta_E(v_x) \cup \Delta_{E^{v_x}}(v \ominus v_x).$$

By point 4 of Theorem 4.7 $\Delta_E(v_x) = \Delta_x(v_x)$. Since x , as an event structure, is the only nonempty stopping prefix of itself, it is clear that $\Delta_x(v_x) = \{x\}$. Hence we have $\Delta_E(v) \supseteq \{x\}$. Symmetrically, we also obtain that $\Delta_E(v) \supseteq \{y\}$. But $v = (v \cap x) \cup (v \cap y)$, so we are done: $\Delta_E(v) = \{x, y\}$.

More generally, we retain that if v has the form $v = \bigcup_{x \in \xi} v_x$, where ξ is a set of initial stopping prefixes, and v_x is maximal in x for every $x \in \xi$, then the covering $\Delta_E(v)$ is given by $\Delta_E(v) = \xi$. Although quite intuitive, this result is not obvious from the only definition of the covering map.

As another example, consider the event structure E depicted in Figure 8. We shall determine the coverings of the two maximal configurations $\omega_1 = (ad)$ and $\omega_2 = (bce)$ of E .

Let $x = \{a, b\}$ be the unique initial stopping prefix of E . Then $\Delta_E(\omega_1)$ and $\Delta_E(\omega_2)$ both contain x . Figure 8 (2) and (3) respectively depict the futures $E^{(b)}$ and

$E^{(a)}$ of configurations (b) and (a), with the associated initial stopping prefixes:

$$\delta_E(b) = \{z, z'\}, \quad \delta_E(a) = \{y\},$$

with $z = \{c\}$, $z' = \{e\}$, $y = \{c, d, e\}$. Hence we obtain $\Delta(\omega_1) = \{x, y\}$ and $\Delta(\omega_2) = \{x, z, z'\}$. This example shows that *branching cells of an event structure may overlap*, although branching cells in a same covering $\Delta_E(v)$ shall not overlap as stated by point 3 of Theorem 4.7.

The latter example also shows that stopped configurations are not stable under concatenation. Indeed, configuration (a) is stopped in E , and configuration (c) is stopped in $E^{(a)}$, but their concatenation $(a) \oplus (c) = (ac)$ is not stopped in E , otherwise it would be maximal in E , which is not.

4.4 Normal Decomposition of Maximal Configurations

We shall now give a systematic way to decompose maximal configurations, we call it the *normal decomposition*. This normal decomposition can be seen as a synchronous scheduling of the concurrent events in a same configuration. This result deeply depends on the local finiteness assumption. This normal decomposition is of interest *per se*. It will be instrumental in proving the Law of Large Numbers.

For E a (locally finite) event structure, we define the *max-min stopping prefix* of E as the upper bound of initial stopping prefixes. That is, denoting the max-min stopping prefix of E by $B_0(E)$:

$$B_0(E) =_{[\text{def}]} \bigcup_{x \in \delta_E(\emptyset)} x,$$

where x ranges over the set of initial stopping prefixes of E . Observe that, since initial stopping prefixes are finitely many according to Theorem 4.7, point 2b, and since each branching cell is finite by according to Theorem 4.7, point 1, the max-min stopping prefix $B_0(E)$ is itself a finite stopping prefix.

We define, for ω a maximal configuration of E , the *normal decomposition* of ω as the following pair of sequences, $(V_n(\omega))_{n \geq 0}$ and $(Z_n(\omega))_{n \geq 1}$:

$$V_0 = \emptyset, \quad n > 0, \quad \begin{cases} Z_n(\omega) = \omega \cap B_0(E^{V_{n-1}(\omega)}) \\ V_n(\omega) = V_{n-1}(\omega) \oplus Z_n(\omega) \end{cases} \quad (11)$$

Theorem 4.8 ([6]) *For every $\omega \in \Omega$, the sequence $(V_n(\omega))_{n \geq 0}$ is a nondecreasing sequence of finite R -stopped configurations satisfying: $\sup_n V_n(\omega) = \omega$.*

See an interpretation of Theorem 4.8 by means of σ -algebras and probability in §8.1, Lemma 8.1.

5 Application of Branching Cells to Probabilistic Event Structures

In this section we apply the previous results on branching cells to construct probabilistic event structures, i.e., to construct a class of probability measures (called distributed

probabilities) on the space of maximal configurations of an event structure. This construction follows the results from [9, 6].

Then we discuss *a posteriori* in § 5.4 the meaning of this construction, as well as the relationship between concurrency and probability is discussed. Also in § 5.4, more general constructions are considered, letting appear the branching cell construction as the *minimal decomposition*, from which coarser constructions can be built up.

5.1 Background on Probability Spaces

We first recall some basic notions from Probability theory, see for example [10].

Measurable Spaces and Measurable Mappings. Let U be a nonempty set. A σ -algebra on U is a collection \mathfrak{F} of subsets of U such that $\emptyset \in \mathfrak{F}$, and \mathfrak{F} is closed under complement and under countable union. The pair (U, \mathfrak{F}) is said to be a measurable space. The sets $A \in \mathfrak{F}$ are called \mathfrak{F} -measurable, or simply the measurable subsets of U . If (U, \mathfrak{F}) and (V, \mathfrak{G}) are two measurable spaces, a mapping $\varphi : U \rightarrow V$ is said to be a *measurable mapping* if $\varphi^{-1}(A) \in \mathfrak{F}$ for any $A \in \mathfrak{G}$. We usually adopt the folklore of probability, where measurable mappings are called *random variables*. A bijective mapping $\varphi : U \rightarrow V$ is said to be *bi-measurable* if both φ and φ^{-1} are measurable.

Let U be a nonempty set. For any collection \mathcal{F} of subsets of U , there is a smallest σ -algebra \mathfrak{F} that contains \mathcal{F} . \mathfrak{F} is called the σ -algebra *generated* by \mathcal{F} . Except if otherwise specified, any finite set U is equipped with its *discrete σ -algebra*, that is \mathfrak{F} is simply the powerset of U .

If (U, \mathfrak{F}) is a measurable space, any measurable subset A is equipped with the σ -algebra \mathfrak{F}^A *induced* by \mathfrak{F} , defined by:

$$\mathfrak{F}^A = \{B \in \mathfrak{F} : B \subseteq A\} = \{B \cap A, B \in \mathfrak{F}\}.$$

Probability Measures. If (U, \mathfrak{F}) is a measurable space, the triple $(U, \mathfrak{F}, \mathbb{P})$ is said to be a *probability space* if \mathbb{P} is a nonnegative set function $\mathbb{P} : \mathfrak{F} \rightarrow \mathbb{R}$ such that $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(\Omega) = 1$, and for any sequence $(A_n)_{n \geq 0}$ of pairwise disjoint measurable subsets, we have $\mathbb{P}(\bigcup_{n \geq 0} A_n) = \sum_{n \geq 0} \mathbb{P}(A_n)$. \mathbb{P} is called a *probability measure*, or simply a *probability*.

If U is a finite set, equipped with the discrete σ -algebra \mathfrak{F} , a probability \mathbb{P} is entirely determined by the values of \mathbb{P} on the singletons $\mathbb{P}(\{x\})$, $x \in U$. We simply note $\mathbb{P}(x) = \mathbb{P}(\{x\})$, and we have $\sum_{x \in U} \mathbb{P}(x) = 1$. Conversely, for any nonnegative function $f : U \rightarrow \mathbb{R}$ such that $\sum_{x \in U} f(x) = 1$, there is a unique probability \mathbb{P} on (U, \mathfrak{F}) such that $\mathbb{P}(x) = f(x)$, defined by $\mathbb{P}(A) = \sum_{x \in A} f(x)$ for $A \subseteq U$.

If $\varphi : U \rightarrow V$ is a measurable mapping from (U, \mathfrak{F}) to (V, \mathfrak{G}) , and if \mathbb{P} is a probability measure on (U, \mathfrak{F}) , the following formula defines \mathbb{Q} as a probability measure on (V, \mathfrak{G}) : $\mathbb{Q}(A) = \mathbb{P}(\varphi^{-1}(A))$ for $A \in \mathfrak{G}$. \mathbb{Q} is called the *image of \mathbb{P} under φ* , and we denote it $\mathbb{Q} = \varphi\mathbb{P}$.

Let $(U, \mathfrak{F}, \mathbb{P})$ be a probability space, and let A be a measurable subset of U such that $\mathbb{P}(A) > 0$. Define the *conditional probability given A* by:

$$\forall B \in \mathfrak{F}^A, \quad \mathbb{P}(B|A) = \frac{\mathbb{P}(B)}{\mathbb{P}(A)}.$$

If $(U, \mathfrak{F}, \mathbb{P})$ is a probability space, still following the usual terminology, we define the (*mathematical*) *expectation* of a real-valued nonnegative random variable f as its integral, denoted by $\mathbb{E}(f)$, so that $\mathbb{E}(f) = \int_{\Omega} f(\omega) d\mathbb{P}(\omega)$. If \mathfrak{G} is a sub- σ -algebra of \mathfrak{F} , for every nonnegative \mathfrak{F} -measurable real-valued function f , there exists a \mathfrak{G} -measurable function g such that $\mathbb{E}(fh) = \mathbb{E}(gh)$ for every nonnegative \mathfrak{G} -measurable real-valued function h . Function g is unique up to a set of probability zero, it is called the *conditional expectation of f given \mathfrak{G}* , denoted by $\mathbb{E}(f|\mathfrak{G})$.

Isomorphisms of Probability Spaces. Let $(U, \mathfrak{F}, \mathbb{P})$ and $(V, \mathfrak{G}, \mathbb{Q})$ be two probability spaces. We say that they are *isomorphic* if there are two measurable subsets $U' \subseteq U$ and $V' \subseteq V$ such that $\mathbb{P}(U') = 1$ and $\mathbb{Q}(V') = 1$, and a bi-measurable bijective mapping $\varphi : U' \rightarrow V'$ such that:

$$\varphi\mathbb{P} = \mathbb{Q}, \quad \varphi^{-1}\mathbb{Q} = \mathbb{P}.$$

In the above equations, \mathbb{P} and \mathbb{Q} must be understood as their restriction to the induced σ -algebras $\mathfrak{F}^{U'}$ and $\mathfrak{G}^{V'}$ respectively.

5.2 Probabilistic Event Structures

An event structure E naturally defines a measurable space as follows. Consider first, for any configuration v of E , the following nonempty subset of Ω_E :

$$\mathcal{S}(v) =_{[\text{def}]} \{\omega \in \Omega_E : \omega \supseteq v\}.$$

$\mathcal{S}(v)$ is called the *shadow* of v . We define the *Borel σ -algebra of Ω_E* as the σ -algebra generated by the collection of shadows $\mathcal{S}(v)$, where v ranges over the finite configurations of E . This is indeed the Borel σ -algebra generated by the restricted Scott topology on Ω_E (see [9, Ch.2, III-1.1] for details). Unless otherwise specified, Ω_E is *always equipped with the Borel σ -algebra*, and thus we simply omit it.

The following definition has already been considered in [11, 12, 9], see also the *probabilistic runs* of [13].

Definition 5.1 (probabilistic event structure, likelihood) A probabilistic event structure is a pair (E, \mathbb{P}) where E is an event structure and \mathbb{P} is a probability measure on the space Ω_E of maximal configurations of E .

If (E, \mathbb{P}) is a probabilistic event structure, we define the *likelihood of \mathbb{P}* as the real-valued function $p : \mathcal{V}_E \rightarrow \mathbb{R}$ defined by:

$$\forall v \in \mathcal{V}_E, \quad p(v) = \mathbb{P}(\mathcal{S}(v)).$$

We say that a probabilistic event structure (E, \mathbb{P}) is *positive* if we have:

$$\forall v \in \mathcal{V}_E, \quad p(v) > 0.$$

Next, consider a probabilistic event structure (E, \mathbb{P}) and define the following two notions:

1. **Restriction to stopping prefixes.** Let B be a stopping prefix of E , and let $\pi_B : \Omega_E \rightarrow \Omega_B$ be the mapping defined by $\pi_B(\omega) = \omega \cap B$, which is well defined according to Lemma 4.3. Mapping π_B is measurable. We define the probability \mathbb{P}_B on Ω_B by:

$$\mathbb{P}_B = \pi_B \mathbb{P},$$

image of \mathbb{P} under π_B . In particular, if B is a finite stopping prefix, (Ω_B, \mathbb{P}_B) is a finite probability space, and we have:

$$\forall v \in \Omega_B, \quad \mathbb{P}_B(v) = \mathbb{P}(\mathcal{S}(v)).$$

2. **Probabilistic future.** Let v be a finite configuration of E such that $p(v) > 0$. $\mathcal{S}(v)$ is then a measurable subset of Ω with positive probability. It is thus equipped with the conditional probability, denoted by \mathbb{P}^v , and defined as follows, for A a measurable subset of $\mathcal{S}(v)$:

$$\mathbb{P}^v(A) = \frac{1}{p(v)} \mathbb{P}(A).$$

Consider the bijective and bi-measurable mapping:

$$\Phi^v : \mathcal{S}(v) \rightarrow \Omega_{E^v}, \quad \omega \mapsto \omega \oplus v.$$

We still denote by \mathbb{P}^v the probability on Ω_{E^v} , image of \mathbb{P}^v under Φ^v . For every finite configuration v with $p(v) > 0$, we define the probabilistic event structure (E^v, \mathbb{P}^v) thus obtained as the *probabilistic future* of v . The likelihood p^v of (E^v, \mathbb{P}^v) is given by:

$$\forall w \in \mathcal{V}_{E^v}, \quad p^v(w) = \frac{1}{p(v)} p(v \oplus w). \quad (12)$$

5.3 Locally Randomized Event Structures

We have shown in [6] that a probabilistic event structure can be naturally defined from the new notion of *locally randomized* event structure. We recall this construction.

Definition 5.2 (locally randomized event structure) A locally randomized event structure is a pair $(E, (p_x)_{x \in X})$, where X is the set of branching cells of E , and for each $x \in X$, p_x is a probability on Ω_x . Say that $(E, (p_x)_{x \in X})$ is positive if

$$\forall x \in X_E, \quad \forall z \in \Omega_x, \quad p_x(z) > 0.$$

Observe that, in this definition, the probabilities p_x are finite probabilities since branching cells are finite by Theorem 4.7, point 1, and thus every Ω_x is *a fortiori* finite. Let $(E, (p_x)_{x \in X})$ be a locally randomized event structure. For B a finite stopping prefix of E , we set:

$$\forall \omega_B \in \Omega_B, \quad \mathbb{Q}_B(\omega_B) = \prod_{x \in \Delta(\omega_B)} p_x(\omega_B \cap x), \quad (13)$$

which is well defined since, according to Theorem 4.7, point 3, $\omega_B \cap x \in \Omega_x$. Remark that, if $B = B_0(E)$ is the max-min stopping prefix of E , then \mathbb{Q}_B coincides with the product probability on Ω_B :

$$\mathbb{Q}_B = \bigotimes_{x \in \delta_E(\emptyset)} p_x.$$

This is the probabilistic counterpart of decomposition $\Omega_B = \prod_{x \in \delta(\emptyset)} \Omega_x$, stated in Equation (8), point 2 of Theorem 4.7. The product form of probability \mathbb{Q}_B manifests that “local actions” associated to initial stopping prefixes $x \in \delta(\emptyset)$ are independent in the probabilistic sense. The family $(\mathbb{Q}_B)_B$ makes E a probabilistic event structure, as expressed by the following theorem:

Theorem 5.3 (distributed product) *Let $(E, (p_x)_{x \in X})$ be a locally randomized event structure.*

1. **Distributed product and distributed probabilities.** *There exists a unique probabilistic event structure (E, \mathbb{P}) such that $\mathbb{P}_B = \mathbb{Q}_B$ for every finite stopping prefix $B \subseteq E$, where \mathbb{P}_B denotes the restriction $\pi_B \mathbb{P}$ of \mathbb{P} to Ω_B . The probability \mathbb{P} is called the distributed product of the family $(p_x)_{x \in X}$, written:*

$$\mathbb{P} = \bigotimes_{x \in X}^{\text{dist}} p_x.$$

For every finite R -stopped configuration v , we have:

$$p(v) = \prod_{x \in \Delta(v)} p_x(v \cap x). \quad (14)$$

In particular, the locally randomized event structure $(E, (p_x)_{x \in X})$ is positive if and only if the probabilistic event structure (E, \mathbb{P}) is positive. Probabilities \mathbb{P} arising from a distributed product are called distributed probabilities.

2. **Restriction to stopping prefixes.** *For every stopping prefix $B \subseteq E$, the restriction $\mathbb{P}_B = \pi_B \mathbb{P}$ coincides with the distributed product of the family $(p_x)_{x \in X_B}$, which is the restriction of $(p_x)_{x \in X}$ to the branching cells of B :*

$$\mathbb{P}_B = \bigotimes_{x \in X_B}^{\text{dist}} p_x.$$

3. **Probabilistic future.** *Assume that $(E, (p_x)_{x \in X})$ is positive, so that (E, \mathbb{P}) is a positive probabilistic event structure. Then for every finite R -stopped configuration v , the probabilistic future (E^v, \mathbb{P}^v) coincides with the distributed product of the locally randomized event structure $(E^v, (p_x)_{x \in X^v})$, obtained by restricting family $(p_x)_{x \in X}$ to set X^v of all branching cells of E^v :*

$$\forall v \in \mathcal{W}_E, \quad \mathbb{P}^v = \bigotimes_{x \in X^v}^{\text{dist}} p_x.$$

REMARK. Point 2 is almost immediate. We mention it explicitly to underline the symmetry with point 3. Formula (14) extends (13) from stopped configurations to R -stopped configurations. Formula (14) also shows that, for confusion-free event structures, the *valuations with independence* defined in [13] are equivalently defined as functions of the form (14) associated with distributed products.

Corollary 5.4 *Let $(E, (p_x)_{x \in X})$ be a locally randomized event structure, and let (E, \mathbb{P}) be the associated distributed product. For any finite and compatible R -stopped configurations v and w , if we set $u = v \cap w$, we have:*

$$p(v \cup w) = \frac{1}{p(u)} p(v) p(w).$$

In words, two compatible configurations v and w are probabilistically independent conditionally on their common past. This property expresses that “concurrency matches probabilistic independence”, at the grain of branching cells [6]. The fact that branching cells are minimal with this property is discussed in § 5.4 below.

On the other hand, if a distributed probability \mathbb{P} on Ω_E is positive, it is shown in [6] that there is a *unique* locally randomized event structure $(E, (p_x)_{x \in X})$ such that \mathbb{P} is the distributed product of $(p_x)_{x \in X}$.

5.4 Discussion: why branching cells and stopping prefixes?

Now that the reader has a complete account on stopping prefixes, branching cells, and their probabilistic counterpart, it is time to discuss the relevance of our constructions.

Our first task was to construct probabilistic event structures as limits of *finite* probabilistic event structures. Projective limits provide a flexible notion of limit, and measure theoretic extension theorems exist in the literature concerning projective limits of probabilities (Prokhorov theorem, see [14, 15] for its application to event structures). Let \mathcal{P} denote the set of finite prefixes of an event structure E . Since \mathcal{P} is directed, it is tempting to see it as an index set, and to apply extension results to the family of event structures \mathcal{P} , by using the family of spaces $(\Omega_P)_{P \in \mathcal{P}}$ collecting maximal executions of P .

Unfortunately, it is generally not true that $(\Omega_P)_{P \in \mathcal{P}}$ is a projective system having Ω as its projective limit. The natural projective system that \mathcal{P} defines is the system $(\Gamma_P)_{P \in \mathcal{P}}$, where Γ_P is defined, for $P \in \mathcal{P}$, by

$$\Gamma_P =_{[\text{def}]} \{\omega \cap P, \omega \in \Omega\}.$$

Some prefixes P satisfy $\Gamma_P = \Omega_P$ (in general only the inclusion $\Omega_P \subseteq \Gamma_P$ holds). Call *intrinsic* a prefix with this property. We have shown elsewhere [15] that, if every event of E belongs to some *finite* intrinsic prefix, then Ω has a natural representation as the projective limit $\Omega = \varprojlim_{P \in \mathcal{P}} \Gamma_P$. In this case *any* regular probability measure \mathbb{P} on Ω will be expressed as a projective limit $\mathbb{P} = \varprojlim_{P \in \mathcal{P}} \mu_P$, where $(\mu_P)_{P \in \mathcal{P}}$ defines a projective system of probability measures on $(\Gamma_P)_{P \in \mathcal{P}}$ (see [14, 15]). Hence our first condition, namely the existence of enough finite intrinsic prefixes, was mainly due to measure theoretic considerations.

We now come to the relationship between concurrency and probability. Our aim was to explore the possibilities of matching concurrency with probabilistic independence. A first observation is that concurrency cannot be seen as totally disconnected from causality. We ask the reader to refer to Figure 3, (b) of §2 for an illustration. In this example, the two events c and e are structurally correlated, although concurrent: the former one occurs if and only if the other one does. Probability cannot relax a hard constraint that is structural in the model. This enlightens the fact that, if one wishes to randomize a sub-event structure B of E independently of the remaining of E , a minimal requirement is that B must be “closed under structural correlation”. Obviously, events in minimal conflict with events of B , and events in the downward closure of B , are correlated with B . Hence it is natural to require that B is a stopping prefix. Fortunately, stopping prefixes are intrinsic according to the above definition and thus this new request does not contradict the first one we discussed.

The above discussion shows that branching cells appear as natural candidates for a fine (the finest) decomposition of configurations for subsequent randomization. We could have as well considered coarser decompositions, by clustering finitely many branching cells together. We would then attach local probabilities to such clusters, thus allowing for correlation between branching cells belonging to the same cluster. A natural starting point would be to axiomatize the essential properties of such clusters—our above discussion draw the directions for this. We preferred not to develop this more axiomatic approach, however, as it would have brought further technicalities into the subject.

6 Markov Nets

In this section, we apply the previous probabilistic constructions to event structures arising from the unfolding of safe Petri nets. Recall that such an event structure E is labeled by the transitions of the considered net. It is therefore natural to consider the subclass of associated locally randomized event structures such that the labeling $x \mapsto p_x$, of the branching cells x of E by local probability p_x , conforms to the labeling of events. That is, we require that the locally randomized event structures $(E, (p_x)_{x \in X})$ satisfy $p_x = p_{x'}$ whenever x, x' are branching cells isomorphic as labeled event structures (some precautions must be taken to make this meaningful; this is the topic of Lemma 6.2 below and the following comment). This leads to the notion of *Markov net*, a proper generalization of discrete Markov chains to true-concurrent systems (see §6.2 for a detailed discussion of the latter claim).

For the unfolding of a safe Petri net, equivalence classes of branching cells up to isomorphism of labeled event structures are finitely many. We call them *dynamic clusters*. We shall argue that dynamic clusters are an appropriate concept of local state, for Petri nets. The main theorem of this paper, namely the Law of Large Numbers (LLN) for Markov nets, supports this claim. Indeed, we show that the LLN holds and generalizes the LLN for Markov chains, provided that the set of dynamic clusters is taken as the state space.

Local finiteness of the unfolding is assumed. This is a non-trivial restriction on the net, although this class of Petri nets is strictly larger than the classes of free-choice

or confusion-free nets. As was said in the Introduction, key elements of our approach remain valid in a more general setting—see the discussion at the end of § 8.2 for details about this claim.

6.1 Definition and First Properties of Markov Nets

\mathcal{N} generically denotes a safe Petri net. Denote by E the canonical event structure that unfolds \mathcal{N} (as recalled in §3.1). For v a finite configuration of E , we denote by $\gamma(v)$ the marking reached in \mathcal{N} after the action of configuration v . If m is a reachable marking, we denote by \mathcal{N}^m the safe Petri net identical to \mathcal{N} , except that \mathcal{N} has m as initial marking. We also denote by E^m the unfolding of \mathcal{N}^m .

Recall that two T -labeled event structures (E, λ) and (E', λ') , i.e., two event structures equipped with mappings $\lambda : E \rightarrow T$ and $\lambda' : E' \rightarrow T'$ are said to be *isomorphic* if there is a mapping $\Phi : E \rightarrow E'$ such that:

1. $\forall e, e' \in E, \quad e \preceq e' \iff \Phi(e) \preceq \Phi(e') \text{ and } e \# e' \iff \Phi(e) \# \Phi(e')$;
2. $\lambda' = \Phi \circ \lambda$.

It is well known that, if v is a finite configuration of E , there is a *unique* isomorphism of labeled event structures $E^m \rightarrow E^v$ (see a proof in [9, Ch.5, 1-2.5]). This makes the notation E^m coherent with our previous notation E^v for the future of configurations, so that we can write $E^v = E^{\gamma(v)}$. We can also rewrite it as follows:

$$\forall v, v' \in \mathcal{V}_E, \quad \gamma(v) = \gamma(v') \Rightarrow E^v = E^{v'}. \quad (15)$$

Finally, for m a reachable marking, we denote by γ^m the map that is defined on the poset of finite configurations of E^m and such that $\gamma^m(w)$ is the marking reached by w from marking m , for $w \in \mathcal{V}_{E^m}$.

Since the reachable markings are finitely many, the futures $E^v = E^{\gamma(v)}$ are finitely many up to isomorphism of labeled event structures. Since each set of branching cells $\delta_E(v)$ is finite, it follows then from Definition 4.6 that branching cells of E are finitely many, *up to an isomorphism of labeled event structures*.

Definition 6.1 (dynamic cluster) *An isomorphism class of branching cells is called a dynamic cluster of \mathcal{N} . We denote by Σ the (finite) set of dynamic clusters of \mathcal{N} . Dynamic clusters are generically denoted by the boldface symbol \mathbf{s} . The equivalence class of branching cell x is denoted by $\langle x \rangle$.*

REMARK. It is shown in [16] that, if the event structure is confusion-free, branching cells can be interpreted as the events of a new event structure, called *choice structure*. The set of dynamic clusters Σ is then a finite alphabet that labels the choice structure. Under certain conditions, the labeled event structure obtained is actually itself the unfolding of a safe Petri net, called the *choice net*. The interested reader is referred to [16] for further details.

Lemma 6.2 *If x and x' are two isomorphic branching cells of the unfolding E , then the isomorphism $x \rightarrow x'$ is unique.*

Proof. Let v be a finite R -stopped configuration of E such that $x \in \delta_E(v)$. Let m' be the set of places b of \mathcal{N} such that there is an event $e \in x$, minimal in x , and with $bF\lambda(e)$; i.e., b is in the preset of $\lambda(e)$. Consider the subnet \mathcal{N}' of \mathcal{N} with same sets of places and transitions, but with initial marking m' . Then the unfolding F of \mathcal{N}' has a unique initial stopping prefix y , isomorphic to x . By the uniqueness property of unfoldings, the isomorphism $x \rightarrow y$ is unique.

Symmetrically, there is a unique isomorphism $x' \rightarrow y$. Hence the isomorphism $x \rightarrow x'$ is unique, which completes the proof of the lemma. \square

As a consequence, for every dynamic cluster \mathbf{s} and every pair of branching cells $x, x' \in \mathbf{s}$, there exists a *canonical* bijection $\phi_{x,x'} : \Omega_x \rightarrow \Omega_{x'}$, namely, the bijection induced by the unique isomorphism $x \rightarrow x'$. Hence, we can consistently consider the space $\Omega_{\mathbf{s}}$. In particular, we may define a real-valued function $g : \Omega_{\mathbf{s}} \rightarrow \mathbb{R}$, by means of a family of functions $(g_x)_{x \in \mathbf{s}}$, $g_x : \Omega_x \rightarrow \mathbb{R}$, such that $g_x = g_{x'} \circ \phi_{x,x'}$ for all $x, x' \in \mathbf{s}$. This is what is meant in the following definition, for the probability $p_{\mathbf{s}}$ on $\Omega_{\mathbf{s}}$ seen as a function on $\Omega_{\mathbf{s}}$.

Definition 6.3 (Markov net, local transition probabilities) *A Markov net is a pair $(\mathcal{N}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma})$, where \mathcal{N} is a finite safe Petri net with locally finite unfolding, and $p_{\mathbf{s}}$ is a probability on the finite set $\Omega_{\mathbf{s}}$ for every $\mathbf{s} \in \Sigma$. Probability $p_{\mathbf{s}}$ is called the local transition probability attached to $\mathbf{s} \in \Sigma$. We assume moreover the following:*

$$\forall \mathbf{s} \in \Sigma, \quad \forall z \in \Omega_{\mathbf{s}}, \quad p_{\mathbf{s}}(z) > 0.$$

A Markov net $(\mathcal{N}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma})$ induces a locally randomized event structure $(E, (p_x)_{x \in X})$ by setting $p_x = p_{\langle x \rangle}$ for every branching cell $x \in X_E$; in turn, the distributed product \mathbb{P} of the family $(p_x)_{x \in X}$ defines a probabilistic event structure (E, \mathbb{P}) . Furthermore, this probabilistic event structure is *positive*, according to point 3 in Theorem 5.3.

Note that, if net \mathcal{N} is composed of two disjoint, and thus non interacting nets $\mathcal{N} = \mathcal{N}_1 \cup \mathcal{N}_2$, then the two components $\mathcal{N}_i, i \in \{1, 2\}$ are independent in the probabilistic sense, i.e., $\mathbb{P} = \mathbb{P}_1 \otimes \mathbb{P}_2$: once again, “probabilistic independence matches concurrency”.

Theorem 6.4 (Homogeneity) *Let $(\mathcal{N}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma})$ be a Markov net, and let \mathbb{P} be the associated distributed probability on Ω . For v a finite R -stopped configuration of E , we use the notations $\gamma(v)$ and $\mathcal{N}^{\gamma(v)}$ introduced above, and we let $\Sigma^{\gamma(v)}$ denote the set of dynamic clusters of $\mathcal{N}^{\gamma(v)}$.*

For any $v \in \mathcal{W}$, the probabilistic future (E^v, \mathbb{P}^v) is associated with Markov net $(\mathcal{N}^{\gamma(v)}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma^{\gamma(v)}})$. Moreover we have:

$$\forall v, v' \in \mathcal{W}, \quad \gamma(v) = \gamma(v') \Rightarrow \mathbb{P}^v = \mathbb{P}^{v'}. \quad (16)$$

Proof. According to Theorem 5.3, point 3, we have, for $v \in \mathcal{W}_E$:

$$\mathbb{P} = \bigotimes_{x \in X}^{\text{dist}} p_x, \quad \mathbb{P}^v = \bigotimes_{x \in X^v}^{\text{dist}} p_x,$$

where X^v is the set of branching cells of E^v . Therefore, \mathbb{P}^v is indeed the distributed product associated with Markov net $(\mathcal{N}^{\gamma(v)}, (p_s)_{s \in \Sigma^{\gamma(v)}})$.

For $v, v' \in \mathcal{W}$ with $\gamma(v) = \gamma(v')$, let $\varphi : E^v \rightarrow E^{v'}$ denote the unique isomorphism of labeled event structures. We denote by $\Phi : X^v \rightarrow X^{v'}$ the induced bijection between the two sets of branching cells. We have:

$$\mathbb{P}^v = \bigotimes_{x \in X^v}^{\text{dist}} p_x = \bigotimes_{y \in X^{v'}}^{\text{dist}} p_{\Phi^{-1}(y)}, \quad (17)$$

by the change of variable $x = \Phi^{-1}(y)$, which is allowed since Φ is bijective. For every branching cell $x \in X^v$, the restriction $\varphi|_x : x \rightarrow \Phi(x)$, with $x \in X^v$, does not change the class of branching cell. Hence, up to isomorphism of branching cells, $\varphi|_x$ is the identity map. Therefore $p_{\Phi^{-1}(y)} = p_y$ for every $y \in X^{v'}$. Hence, from (17), we get:

$$\mathbb{P}^v = \bigotimes_{y \in X^{v'}}^{\text{dist}} p_y = \mathbb{P}^{v'}.$$

This completes the proof of the theorem. \square

Equation (16) expresses that the memory of Markov nets is entirely summarized by the current marking: the probabilistic future of a $v \in \mathcal{W}$ only depends on the final marking $\gamma(v)$, and not on the entire history v . It is the probabilistic counterpart of Equation (15). In the setting of Markov chains, this is equivalent to the time-invariance property of the transition matrix, characteristic of *homogeneous Markov chains* (see e.g., [10, 7]). An important consequence of homogeneity, for Markov chains, is the so-called *strong Markov property*; see a generalization for Markov nets in [17].

It will be convenient to use the following terminology and notation:

Definition 6.5 (recursively stopped marking) *We say that a marking m is recursively stopped if there is a finite recursively stopped configuration v such that $\gamma(v) = m$. The set of recursively stopped markings is denoted by \mathbf{M}_{rs} .*

Thanks to property (16), we may define for each $m \in \mathbf{M}_{\text{rs}}$, the probabilistic event structure (E^m, \mathbb{P}^m) , by defining \mathbb{P}^m as the probability \mathbb{P}^v on Ω_{E^m} , whenever v is a finite R -stopped configuration such that $\gamma(v) = m$. Moreover, Theorem 6.4 says that \mathbb{P}^m is the probability associated with the Markov net $(\mathcal{N}^m, (p_s)_{s \in \Sigma^m})$.

6.2 Markov Nets as a Generalization of Markov Chains

For further referencing it will be useful to see how Markov nets are a generalization of Markov chains. To this end, we construct for each Markov chain a canonical Markov net that simulates the chain.

Let $(X_n)_{n \geq 0}$ be a Markov chain defined on a finite state space P , with transition matrix $(\rho_{i,j})_{(i,j) \in P \times P}$ and initial state s_0 . Consider the following safe Petri net $\mathcal{N} = (P, T, F, m_0)$. The set of places P coincides with the set of states of the chain, and $m_0 = \{s_0\}$. The set T of transitions is defined by this rule: $T \subseteq P \times P$, with:

$$\forall (i, j) \in P \times P, \quad (i, j) \in T \iff \rho_{i,j} > 0.$$

Then the flow relation $F \subseteq (P \times T) \cup (T \times P)$ is naturally defined by:

$$\forall (i, j) \in P \times P, \quad (i, j) \in T \Rightarrow i F (i, j) \quad \text{and} \quad (i, j) F j.$$

Each reachable marking of \mathcal{N} is a singleton, and there is a one-to-one and onto correspondence between paths of the chain and firing sequences of the net. The unfolding of \mathcal{N} is a tree. It is thus clear that the set Σ of dynamic clusters of \mathcal{N} is given by $\Sigma = \{s_i, i \in P\}$, with:

$$\forall i \in P, \quad s_i = \{(i, j), j \in P \text{ such that } (i, j) \in T\}.$$

Hence, for Markov chains, *dynamic clusters identify with states* of the chain.

To define a Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$, it remains only to define the family of local transition probabilities $(p_s)_{s \in \Sigma}$. This is done as follows, using the fact that maximal configurations of a cluster s_i identify with the transitions (i, j) of s_i :

$$\forall i \in P, \quad \forall j \in P, \quad (i, j) \in s_i \Rightarrow p_{s_i}(i, j) = \rho_{i,j}.$$

That is, *local transition probabilities are given by the rows of the transition matrix*. Remark that we have by construction $p_s(z) > 0$, for all $s \in \Sigma$ and $z \in \Omega_s$.

Finally, we must prove that, if $(\Omega, \mathfrak{F}, \mathbb{P})$ is the probability space associated with the executions of Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$, and if $(\Xi, \mathfrak{G}, \mathbb{Q})$ is the probability space associated with the executions of the Markov chain $(X_n)_{n \geq 0}$, there is an isomorphism of probability spaces:

$$(\Omega, \mathfrak{F}, \mathbb{P}) \rightarrow (\Xi, \mathfrak{G}, \mathbb{Q}).$$

This can be done by hand, by checking that formula (14) that defines the likelihood function for the Markov net coincides with the equivalent for the Markov chain. But this is also a consequence of the more general result stated in Lemma 8.1 below, §8.1.

6.3 Recurrent Nets

Recurrent Markov chains are chains in which almost every trajectory returns infinitely often to the initial state. As a generalization, recurrent Markov nets will be nets in which almost every firing sequence returns infinitely often to the initial marking. The above statement calls for distinguishing between two types of returns: “global” returns, where all tokens must leave the initial marking before returning to it, and “local” returns, where some “minimal” amount of tokens move. Global return is investigated next; the study of local return is postponed to §8.4.

Definition 6.6 (return operator) Let m_0 be the initial marking of \mathcal{N} . We define the return operator R on Ω as follows: for every $\omega \in \Omega$,

$$R(\omega) = \min\{v \in \mathcal{W}_E : v \subseteq \omega, \gamma(v) = m_0, \text{Min}(E) \cap \text{Min}(E^v) = \emptyset\}, \quad (18)$$

with the convention that $R(\omega) = \omega$ if the set of such v is empty.

The condition $\text{Min}(E) \cap \text{Min}(E^v) = \emptyset$ says that configuration v has moved all the tokens in the net. In other words, we do not allow the return to leave some part of the marking untouched, while acting on the other tokens.

It is well known that, if v, v' are two compatible configurations such that $\gamma(v) = \gamma(v') = m_0$, then $\gamma(v \cap v') = m_0$. Furthermore, the intersection of compatible R -stopped configurations is still R -stopped thanks to Theorem 4.7, point 5. Finally, it is also readily checked that the intersection of configurations satisfying the condition $\text{Min}(E) \cap \text{Min}(E^v) = \emptyset$ also satisfies this condition. Therefore, if the set of configurations in the right member of (18) is nonempty, the minimum is well defined and is finite.

The study of properties of R is postponed to §8.3. For the moment, we define the successive returns by:

$$R_0 = \emptyset, \quad \forall n > 0, \quad R_n(\omega) = R_{n-1}(\omega) \oplus R(\omega \setminus R_{n-1}(\omega)),$$

with the convention that $R_n(\omega) = \omega$ if $R_{n-1}(\omega) = \omega$. We mention without proof that the successive returns $(R_n)_{n \geq 0}$ are all finite with probability either 0 or 1 [17], which extends to nets a classical result for Markov chains.

Definition 6.7 (recurrent nets) We say that a Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$ is recurrent if the successive returns satisfy $R_n(\omega) \neq \omega$ for all $n \geq 0$, with probability 1.

If the considered net is a simulated Markov chain, as described in §6.2, then this definition of recurrence reduces to the classical notion of recurrence for Markov chains [10, 7], and the R_n are the n^{th} returns to the initial state of the chain. Also, the following lemma extends a classical result for Markov chains [17]:

Lemma 6.8 If $(\mathcal{N}, (p_s)_{s \in \Sigma})$ is a recurrent Markov net, then $(\mathcal{N}^m, (p_s)_{s \in \Sigma^m})$ is recurrent for every $m \in \mathbf{M}_{\text{rs}}$, and $\Sigma^m = \Sigma$.

7 The Law of Large Numbers (LLN)

For our study of the Law of Large Numbers, we focus on recurrent nets.

7.1 What is the proper notion of LLN, for Markov nets?

For finite recurrent Markov chains, the LLN states as follows. Let Σ be the finite state space of a Markov chain $(X_k)_{k \geq 1}$, and let $f : \Sigma \rightarrow \mathbb{R}$ be a test function. The sums $S_n(f) = \sum_{k=1}^n f(X_k)$ are called *ergodic sums*, and the LLN studies the limit, for $n \rightarrow \infty$, of the *ergodic means*: $M_n(f) = \frac{1}{n} S_n(f)$. It is always true that $M_n(f)$

has a limit when n goes to infinity, and this limit is constant with probability 1 (i.e., independent of the random outcome) under some mild conditions. The limit is then given by $\int_{\Sigma} f(x) d\mu(x)$, where μ is a probability measure on Σ called the *stationary measure* of the chain.

Hence, what the LLN says, for Markov chains, is that the temporal means $M_n(f)$ have a limit, equal to the spatial mean given by the expectation of f w.r.t. the stationary measure. The coefficients $\mu(x)$, for $x \in \Sigma$, appear as the *asymptotic presence rate* of the chain in state x .

In extending the LLN to Markov net \mathcal{N} , we are faced with two difficulties:

1. What is the proper concept of state? What are the associated ergodic sums?
2. What replaces counter n , since time is not totally ordered?

Corresponding answers are:

1. The set Σ of dynamic clusters of \mathcal{N} is taken as the state space—see the discussion of §6.2.
2. For v a R -stopped configuration, the number of branching cells contained in the covering $\Delta(v)$ is taken as the “duration” of v .

Since we consider dynamic clusters as our state space, test functions are simply functions defined on Σ :

Definition 7.1 (state function) *If Σ denotes the set of dynamic clusters of a net \mathcal{N} , we call state function any real-valued function $f : \Sigma \rightarrow \mathbb{R}$.*

State functions form a vector space of finite dimension ($=\text{Card}(\Sigma)$). The *concurrent ergodic sums* associated with a state function $f : \Sigma \rightarrow \mathbb{R}$ are defined as the function $\langle f, \cdot \rangle$:

$$\langle f, \cdot \rangle : \mathcal{W} \rightarrow \mathbb{R}, \quad \forall v \in \mathcal{W}, \quad \langle f, v \rangle = \sum_{x \in \Delta(v)} f(\langle x \rangle),$$

where we recall that $\langle x \rangle$ denotes the dynamic cluster defined by the class of branching cell x . The scalar product notation is justified since $\langle f, v \rangle$ is linear w.r.t. its left argument for the usual addition of functions, and additive w.r.t. its right argument for the concatenation \oplus of configurations.

EXAMPLE. Let $\mathbf{1}$ be the unit state function, defined by:

$$\forall s \in \Sigma, \quad \mathbf{1}(s) = 1. \tag{19}$$

Then $\langle \mathbf{1}, v \rangle$ counts the number of branching cells contained in $\Delta(v)$. This example will be of repeated use in the sequel.

The *concurrent ergodic means* $M(f, \cdot) : \mathcal{W} \rightarrow \mathbb{R}$ associated with a state function f are defined as the following ratios:

$$\forall v \in \mathcal{W}, \quad M(f, v) = \frac{1}{\langle \mathbf{1}, v \rangle} \langle f, v \rangle.$$

The LLN is concerned by the limit:

$$\lim_{v \subseteq \omega, v \rightarrow \omega} M(f, v),$$

where this limit is meant in a sense we shall make precise. The following notion of *stopping operator* will be central in this respect. Stopping operators generalize for concurrent systems the classical notion of *stopping times* [7, 10] for sequential stochastic processes in discrete time; see [17] for a detailed discussion.

Definition 7.2 (stopping operator, regular sequences) *A measurable mapping $V : \Omega \rightarrow \mathcal{W}$, satisfying $V(\omega) \subseteq \omega$ for all $\omega \in \Omega$, is called a stopping operator if for all $\omega, \omega' \in \Omega$, we have:*

$$\forall \omega, \omega' \in \Omega, \quad \omega' \supseteq V(\omega) \Rightarrow V(\omega') = V(\omega). \quad (20)$$

Say that a sequence $(V_n)_{n \geq 1}$ of stopping operators is regular if the following properties are satisfied:

1. $V_n \subseteq V_{n+1}$ for all $n \geq 1$, and $\bigcup_n V_{n \geq 1}(\omega) = \omega$, with probability 1;
2. There exists a constant $k > 0$ such that:

$$\forall n \geq 1, \quad \langle \mathbf{1}, V_n(\omega) \rangle \geq k \times n, \quad \text{with probability 1,}$$

where $\mathbf{1}$ is the unit state function defined in (19).

EXAMPLE. *Stopping prefixes as stopping operators.* As an example, let B be a stopping prefix of E . Then the map $V_B : \Omega \rightarrow \Omega_B$, $\omega \mapsto \omega \cap B$, is a stopping operator. Indeed, $V_B(\omega)$ is R -stopped, and $V_B(\omega) \subseteq \omega$. Finally, let $\omega, \omega' \in \Omega$ such that $\omega' \supseteq V_B(\omega)$. Then $\omega' \cap B \supseteq \omega \cap B$. But, since $\omega \cap B$ is maximal in B , it implies that $\omega' \cap B = \omega \cap B$, which is exactly $V_B(\omega') = V_B(\omega)$. Hence V_B is indeed a stopping operator, as announced.

Not all stopping operators may be represented by this way. However, if the net arises from a Markov chain as in §6.2, then every stopping operator can be represented by a stopping prefix.

We use stopping operators to express the notion of convergence as follows:

Definition 7.3 (convergence of ergodic means) *For f a state function, we say that the ergodic means $M(f, \cdot)$ converge to a function $\mu : \Omega \rightarrow \mathbb{R}$ if for every regular sequence $(V_n)_{n \geq 1}$ of stopping operators,*

$$\lim_{n \rightarrow \infty} M(f, V_n(\omega)) = \mu(\omega), \quad \text{with probability 1.} \quad (21)$$

The important point of this definition is that the limit μ does not depend on the regular sequence of stopping operators. It is thus intrinsic to f and \mathbb{P} .

REMARK. *The need for synchronization.* Concurrency prevents (21) from holding for general recurrent Markov nets, as the following example shows. Assume that \mathcal{N}

decomposes as a disjoint union $\mathcal{N} = \mathcal{N}^1 \cup \mathcal{N}^2$, i.e., the two components \mathcal{N}^1 and \mathcal{N}^2 do not interact at all. The unfolding E of net \mathcal{N} is the union of the unfoldings E^1 and E^2 of nets \mathcal{N}^1 and \mathcal{N}^2 respectively, and the set of dynamic clusters Σ of \mathcal{N} is the union of the sets Σ^1 and Σ^2 of dynamic clusters of \mathcal{N}^1 and \mathcal{N}^2 . We assume that both nets \mathcal{N}^1 and \mathcal{N}^2 are recurrent. Let f be the state function defined by:

$$\forall \mathbf{s} \in \Sigma, \quad f(\mathbf{s}) = \begin{cases} 1, & \text{if } \mathbf{s} \in \Sigma^1, \\ 0, & \text{if } \mathbf{s} \in \Sigma^2. \end{cases}$$

The associated ergodic sums $\langle f, v \rangle$ count the number of occurrences of dynamic clusters from Σ^1 in v . As shown by Lemma 7.4 below, there are two regular sequences $V^i = (V_n^i)_{n \geq 0}$, of stopping operators of \mathcal{N}^i respectively, for $i = 1, 2$, such that for some constants $k, K > 0$:

$$i = 1, 2, \quad \forall n > 1, \quad k \times n \leq \langle \mathbf{1}, V_n^i \rangle \leq K \times n.$$

Define, for $n > 1$, $V_n = V_n^1 \cup V_n^2$, and define symmetrically $W_n = V_n^1 \cup V_n^2$. Then $(V_n)_{n > 1}$ and $(W_n)_{n > 1}$ are two regular sequences of stopping operators of E . We compute the ergodic means of f along each of the sequences $(V_n)_{n > 1}$ and $(W_n)_{n > 1}$. We have:

$$\begin{aligned} M(f, V_n) &= M(f, V_n^1 \cup V_n^2) = \frac{\langle \mathbf{1}, V_n^1 \rangle}{\langle \mathbf{1}, V_n^1 \rangle + \langle \mathbf{1}, V_n^2 \rangle} \\ &\geq \frac{kn^2}{Kn^2 + Kn} \xrightarrow{n \rightarrow \infty} \frac{k}{K} > 0. \end{aligned}$$

On the other hand:

$$\begin{aligned} M(f, W_n) &= M(f, V_n^1 \cup V_n^2) = \frac{\langle \mathbf{1}, V_n^1 \rangle}{\langle \mathbf{1}, V_n^1 \rangle + \langle \mathbf{1}, V_n^2 \rangle} \\ &\leq \frac{Kn}{kn + kn^2} \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

Hence, the limit of the ergodic means depends on the particular regular sequence of stopping operators chosen. Thus, the convergence of ergodic means does not hold in the sense of Definition 7.3.

Clearly, the concurrency properties of \mathcal{N} is the very cause of this difficulty in the example; more precisely, the total absence of synchronization brings this behavior. We show below “how much synchronization” is needed for the LLN to be valid.

The existence of regular sequences of stopping operators results from the following lemma. Recall that we have defined in §4.4 the normal decomposition $(V_n(\omega), Z_{n+1}(\omega))_{n \geq 0}$ of a maximal configuration ω . Since V_n and Z_n are (measurable) mappings defined on Ω , we now see this sequence as a sequence of operators, as in Definition 7.2.

Lemma 7.4 For every $n \geq 0$, the mapping $\omega \mapsto V_n(\omega)$ resulting from the normal decomposition of maximal configuration ω is a stopping operator. There is a constant $K > 0$ such that:

$$\forall n \geq 0, \quad \langle \mathbf{1}, V_n \rangle \leq K \times n. \quad (22)$$

If the net is recurrent, then $(V_n)_{n \geq 0}$ is a regular sequence of stopping operators.

Proof. By construction, $V_n(\omega) \subseteq \omega$ for all $\omega \in \Omega$. We prove the other point of the definition of stopping operators, stated by Equation (20) in Definition 7.2, by induction on $n \geq 0$. This is trivial for $n = 0$; assume that Equation (20) holds for V_n until $n \geq 0$. Let $\omega, \omega' \in \Omega$ such that $\omega' \supseteq V_{n+1}(\omega)$. Then in particular $\omega' \supseteq V_n(\omega)$, so that $V_n(\omega') = V_n(\omega)$ by the induction hypothesis. Put $v = V_n(\omega)$, $\xi = \omega \ominus v$ and $\xi' = \omega' \ominus v$, so that $\xi, \xi' \in \Omega_{E^v}$. We have:

$$Z_{n+1}(\omega') = \xi' \cap B_0(E^v) \supseteq \xi \cap B_0(E^v).$$

Since $\cdot \cap B_0(E^v)$ is a stopping operator of Ω_{E^v} , since it is defined by means of the stopping prefix $B_0(E^v)$, it follows that $Z_{n+1}(\omega') = \xi' \cap B_0(E^v) = \xi \cap B_0(E^v) = Z_{n+1}(\omega)$. Finally, $V_{n+1}(\omega') = V_n(\omega') \oplus Z_{n+1}(\omega') = V_{n+1}(\omega)$, which completes the proof of (20). This shows that every V_n is a stopping operator.

We now show Equation (22). According to Equation (10) in Theorem 5, we have:

$$\forall n > 0, \quad \Delta(V_n) = \Delta(V_{n-1}) \cup \Delta(Z_n), \quad \Delta(V_{n-1}) \cap \Delta(Z_n) = \emptyset,$$

with Δ the covering map. Therefore:

$$\forall n \geq 0, \quad \langle \mathbf{1}, V_n \rangle = \text{Card}(\Delta(V_n)) = \sum_{i=1}^n \text{Card}(\Delta(Z_i)). \quad (23)$$

As observed in §4.3, since Z_i has the form $Z_i = \bigcup_{x \in \delta_E(V_{i-1})} Z_i \cap x$ with $Z_i \cap x \in \Omega_x$ for $x \in \delta_E(V_{i-1})$, the covering $\Delta(Z_i)$ is given by $\Delta(Z_i) = \delta_E(V_{i-1})$. Therefore, according to point 2b of Theorem 4.7, there is a constant K such that $\text{Card}(\Delta(Z_i)) \leq K$. Hence it follows from (23) that:

$$\forall n \geq 0, \quad \langle \mathbf{1}, V_n \rangle \leq K \times n,$$

which is (22).

Now we show that, if \mathcal{N} is recurrent, $(V_n)_{n \geq 0}$ is a regular sequence of stopping operators. It is clear that $V_n \subseteq V_{n+1}$ for all $n \geq 0$. Moreover, according to Theorem 4.8, $\bigcup_{n \geq 0} V_n(\omega) = \omega$. Hence point 1 in Definition 7.2 is satisfied.

Since \mathcal{N} is recurrent (Definition 6.7), with probability 1, ω is infinite since it contains infinitely many returns to the initial marking. Therefore, for each $i \geq 0$, V_i is not maximal, and thus $B_0(E^{V_i}) \neq \emptyset$. In particular, $\text{Card}(\Delta(Z_{i+1})) \geq 1$. Hence, using Equation (23), we obtain $\langle \mathbf{1}, V_n \rangle \geq n$, for all $n \geq 0$. Hence $(V_n)_{n \geq 0}$ satisfies the definition of a regular sequence of stopping operators with $k = 1$. This completes the proof of the lemma. \square

7.2 A Probabilistic Measure of Synchronization and the LLN

If, in an execution $\omega \in \Omega$, we block a token in some place b , we measure the “loss of synchronization” of the system by counting the number of branching cells that can be traversed *without moving the blocked token*. This length defines an integer random variable. A reasonable assumption is that this random variable has finite mean. We detail this definition below.

Let m be a reachable marking of a Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$, with $m \in \mathbf{M}_{rs}$, and let b be a place of m . Denote by $\lambda(e)$ the transition of \mathcal{N} that labels an event $e \in E$. Consider an element $\omega \in \Omega_{E^m}$, and define:

$$K_m(b, \omega) = \sup\{v \in \mathcal{W}_{E^m} : v \subseteq \omega, \forall e \in v, \lambda(e) \notin b^\bullet\},$$

where b^\bullet denotes the postset of b (i.e., the set of transitions t in \mathcal{N} directly after b). $K_m(b, \omega)$ describes the maximal R -stopped sub-configuration that ω allows without using the token in b . Next, define the integer $L_m(b, \omega)$ by:

$$\forall \omega \in \Omega_{E^m}, \quad L_m(b, \omega) = \begin{cases} \infty, & \text{if } K_m(b, \omega) \text{ is infinite,} \\ \langle \mathbf{1}, K_m(b, \omega) \rangle, & \text{otherwise.} \end{cases}$$

Hence $L_m(b, \omega)$ counts the number of branching cells of $K_m(b, \omega)$. It is thus a measure of non-synchronization at ω : the larger $L_m(b, \omega)$, the more sub-processes of ω can progress without synchronizing with b . For each recursively stopped marking m , and each $b \in m$, $L_m(b, \cdot)$ is now an integer random variable $\Omega_{E^m} \rightarrow \mathbb{N} \cup \{\infty\}$. Its integral is thus well-defined, although it may be infinite. We shall thus consider the following definition:

Definition 7.5 (integrable concurrency height) *Say that Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$ has integrable concurrency height if for each $m \in \mathbf{M}_{rs}$, and for each place $b \in m$:*

$$\mathbb{E}^m(L_m(b, \cdot)) < \infty,$$

where \mathbb{E}^m denotes the mathematical expectation under probability \mathbb{P}^m .

This conditions expresses that the expected time of deadlock for a partially blocked system is finite.

EXAMPLES. Clearly, the example given above of a net consisting of the union $\mathcal{N}^1 \cup \mathcal{N}^2$ of two non-interacting and recurrent nets has not integrable concurrency height. In contrast, if a Markov net arises from of Markov chain through the construction of §6.2, then the integrable concurrency height condition is automatically fulfilled, since the variables $L_m(b, \cdot)$ identically vanish. On the other hand, the two nets depicted in Figure 9 have integrable concurrency height.

Theorem 7.6 (Law of Large Numbers) *Let $(\mathcal{N}, (p_s)_{s \in \Sigma})$ be a Markov net, that we assume recurrent and with integrable concurrency height. Then:*

1. *For any state function f , the ergodic means $M(f, \cdot)$ converge in the sense of Definition 7.3 to a function $\mu(f, \cdot) : \Omega \rightarrow \mathbb{R}$.*




Figure 9: Two recurrent Markov nets with integrable concurrency height. Places with same labels (A and B) are identified. In (a), the random variables $L_m(b, \cdot)$ are bounded. This condition is too restrictive in general, as shown by the example (b), where the random variables $L_m(b, \cdot)$ are integrable without being bounded.

2. Except possibly on a set of zero probability, $\mu(f, \cdot) = \mu(f)$ is constant. There exists a family of numbers $(\alpha(\mathbf{s}))_{\mathbf{s} \in \Sigma}$ such that, for every state function $f : \Sigma \rightarrow \mathbb{R}$, the constant $\mu(f)$ is given by:

$$\mu(f) = \sum_{\mathbf{s} \in \Sigma} \alpha(\mathbf{s}) f(\mathbf{s}). \quad (24)$$

3. The coefficients $\alpha(\mathbf{s})$ satisfy:

$$\forall \mathbf{s} \in \Sigma, \quad 0 < \alpha(\mathbf{s}) \leq 1, \quad \text{and} \quad \sum_{\mathbf{s} \in \Sigma} \alpha(\mathbf{s}) = 1.$$

Hence $(\alpha(\mathbf{s}))_{\mathbf{s} \in \Sigma}$ defines a probability on Σ , we call it the stationary measure of the Markov net.

The proof of the theorem is the topic of §8. A close look at the proof shows that it is enough, for Theorem 7.6 to hold, that the probability \mathbb{P} satisfies the homogeneity condition; the stronger property that \mathbb{P} is distributed is not required.

7.3 Interpretation of the Stationary Measure

We give an interpretation of the coefficients $\alpha(\mathbf{s})$ mentioned in Theorem 7.6. For every $\mathbf{s} \in \Sigma$, the coefficient $\alpha(\mathbf{s})$ shall be considered as the asymptotic rate of occurrence of \mathbf{s} in a typical execution $\omega \in \Omega$. Indeed, consider, for $\mathbf{s}_0 \in \Sigma$, the state function $\mathbf{1}_{\mathbf{s}_0}$ defined by:

$$\forall \mathbf{s} \in \Sigma, \quad \mathbf{1}_{\mathbf{s}_0}(\mathbf{s}) = \begin{cases} 1, & \text{if } \mathbf{s} = \mathbf{s}_0, \\ 0, & \text{otherwise.} \end{cases} \quad (25)$$

The ergodic sums $\langle \mathbf{1}_{\mathbf{s}_0}, v \rangle$ count, for $v \in \mathcal{W}$, the number of occurrences of \mathbf{s}_0 in v . Applying the LLN to $\mathbf{1}_{\mathbf{s}_0}$, we get that, for every regular sequence of stopping operators $(V_n)_{n \geq 0}$, and with probability 1:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\text{number of occurrences of } \mathbf{s}_0 \text{ in } V_n}{\text{number of branching cells in } V_n} &= \lim_{n \rightarrow \infty} \frac{\langle \mathbf{1}_{\mathbf{s}_0}, V_n \rangle}{\langle \mathbf{1}, V_n \rangle} \\ &= \lim_{n \rightarrow \infty} M(\mathbf{1}_{\mathbf{s}_0}, V_n) = \mu(\mathbf{1}_{\mathbf{s}_0}) = \alpha(\mathbf{s}_0). \end{aligned} \quad (26)$$

The expression (26) shows that $\alpha(\mathbf{s}_0)$ is the asymptotic occurrence rate of \mathbf{s}_0 in a typical ω , whatever the regular sequence of stopping operators V_n is. This justifies the name of density coefficient.

If the net arises from a recurrent Markov chain through the construction of §6.2, then the integrable concurrency height is satisfied, as already seen, and thus the theorem applies. The stationary measure of the net coincides with the stationary measure of the chain, in the usual sense. This again reveals that dynamic clusters play the role of local states for concurrent systems.

We now examine an example that reveals that the LLN stated in Theorem 7.6 is not fine enough to describe the asymptotic behavior of Markov nets. We will thus need to refine it.

Let $(\mathcal{N}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma})$ be a Markov net, recurrent and with integrable concurrency height as in Theorem 7.6, with unfolding (E, λ) . Pick t and t' two transitions of \mathcal{N} . For v a finite R -stopped configuration of E , let $r_{t,t'}(v)$ be the ratio of occurrences of t and t' in v . That is:

$$r_{t,t'}(v) = \frac{\text{Card}\{e \in v : \lambda(e) = t\}}{\text{Card}\{e \in v : \lambda(e) = t'\}}. \quad (27)$$

We would like to know if this ratio has a limit when v grows to some $\omega \in \Omega_E$. This would tell us how much, asymptotically, transition t fires as compared to t' . However, the ratio $r_{t,t'}(v)$ cannot be expressed as a ratio of the form $\langle f, v \rangle / \langle f', v \rangle$, where f and f' would be state functions. Indeed, the occurrence of transition t and t' in an element $z \in \Omega_{\mathbf{s}}$, with \mathbf{s} a dynamic cluster, depends on z , and not on \mathbf{s} only. Hence state functions are not fine enough to evaluate quantities of the type $r_{t,t'}(v)$. This is the reason why we introduce *extended state functions*.

7.4 Extended state functions and the Extended LLN

We begin with the definition.

Definition 7.7 (extended state function) We call extended state function a finite family $f = (f(\mathbf{s}, \cdot))_{\mathbf{s} \in \Sigma}$ of real-valued functions $f(\mathbf{s}, \cdot) : \Omega_{\mathbf{s}} \rightarrow \mathbb{R}$, where $\mathbf{s} \in \Sigma$.

Extended state functions extend state functions: indeed, a state function $f : \Sigma \rightarrow \mathbb{R}$ is simply an extended state function that is constant on every $\Omega_{\mathbf{s}}$. That is, $f(\mathbf{s}, z) = f(\mathbf{s})$ for all $\mathbf{s} \in \Sigma$ and $z \in \Omega_{\mathbf{s}}$. In particular, we keep the notation $\mathbf{1}$ to denote the unit extended state function, defined by $\mathbf{1}(\mathbf{s}, z) = 1$ for all $\mathbf{s} \in \Sigma$ and $z \in \Omega_{\mathbf{s}}$.

Ergodic sums and means of extended state functions are defined in a way that extends the definition of ergodic sums and means for state functions. The ergodic sum of an extended state function f along a finite R -stopped configuration v is defined by:

$$\langle f, v \rangle = \sum_{x \in \Delta(v)} f(\langle x \rangle, v \cap x),$$

and the ergodic means are defined by:

$$\forall v \in \mathcal{W}_E, \quad M(f, v) = \frac{\langle f, v \rangle}{\langle \mathbf{1}, v \rangle}.$$

If f is an extended state function, we define the convergence of the ergodic means $M(f, \cdot)$ to a function $\mu : \Omega \rightarrow \mathbb{R}$ as in Definition 7.3 for state functions.

EXAMPLE. The ratios of the form $r_{t,t'}(v)$ defined in (27) can be expressed by ergodic means of extended state functions. Indeed, consider the extended state functions N_t and $N_{t'}$ defined by:

$$\forall \mathbf{s} \in \Sigma, \quad \forall z \in \Omega_{\mathbf{s}}, \quad N_t(\mathbf{s}, z) = \text{Card}\{e \in z : \lambda(e) = t\}, \quad (28)$$

and similarly for $N_{t'}$. Then, for v finite and R -stopped, the ratio $r_{t,t'}(v)$ is given by:

$$r_{t,t'}(v) = \frac{\langle N_t, v \rangle}{\langle N_{t'}, v \rangle} = \frac{\langle N_t, v \rangle}{\langle \mathbf{1}, v \rangle} \times \frac{\langle \mathbf{1}, v \rangle}{\langle N_{t'}, v \rangle} = \frac{M(N_t, v)}{M(N_{t'}, v)}. \quad (29)$$

Hence, if we know that ergodic means of extended state functions have a limit, we shall conclude that the ratios $r_{t,t'}(v)$ also have a limit. This is the topic of the following result, the proof of which is postponed in §8.

Theorem 7.8 (extended LLN) *Let $(\mathcal{N}, (p_s)_{s \in \Sigma})$ be a Markov net, that we assume recurrent and with integrable concurrency height. Let $(\alpha(s))_{s \in \Sigma}$ denote the stationary measure of the Markov net.*

Then for every extended state function f , the ergodic means $M(f, \cdot)$ converge to a function $\mu(f, \cdot) : \Omega \rightarrow \mathbb{R}$. Moreover, $\mu(f, \cdot) = \mu(f)$ is constant with probability 1, given by:

$$\mu(f) = \sum_{s \in \Sigma} \alpha(s) p_s(f), \quad \text{with: } p_s(f) = \sum_{z \in \Omega_s} p_s(z) f(s, z), \quad s \in \Sigma. \quad (30)$$

Remark that, in case of an extended state function defined by a state function, the expression (30) giving $\mu(f)$ extends the expression (24) of the first LLN. Indeed, since $f(s, \cdot) = f(s)$ is constant for every s , and since p_s sums up to 1, $p_s(f) = f(s)$ and thus $\mu(f) = \sum_{s \in \Sigma} \alpha(s) f(s)$.

We shall now complete the example of ratios $r_{t,t'}$ defined in (27). Using Equations (29) and (30), we have for every regular sequence $(V_n)_{n \geq 0}$ of stopping operators, and with probability 1:

$$\lim_{n \rightarrow \infty} r_{t,t'}(V_n) = \lim_{n \rightarrow \infty} \frac{M(N_t, V_n)}{M(N_{t'}, V_n)} = \frac{\mu(N_t)}{\mu(N_{t'})},$$

where N_t is defined in (28), and $N_{t'}$ is defined similarly. Hence, the ratios $R_{t,t'}$ have a limit with probability 1, and this limit is the same, with probability 1, for all $\omega \in \Omega$.

REMARK. It is not clear how to extend the LLN, for example, to functions defined on “pairs of successive clusters”. This is done for a Markov chain $(X_n)_{n \geq 0}$ by considering $(X_n, X_{n-1})_{n \geq 1}$, which is again a Markov chain. Our attempts to get a similar construction for nets were not fruitful.

8 Proof of the Law of Large Numbers

The aim of this section is to prove the LLN, as stated in Theorems 7.6 and 7.8. The outline of the proof is as follows:

1. We first fix a particular sequence of stopping operators and examine the convergence of ergodic means along this particular sequence. We show that these ergodic means can be expressed by ergodic means associated with some homogeneous (usual) Markov chain, defined over some special, huge, state space. This shows the convergence of ergodic means to a function $\mu : \Omega \rightarrow \mathbb{R}$ for this particular sequence (Lemma 8.11).

2. We show that the limit obtained remains unchanged if we modify the regular sequence of stopping operators, and that the function μ is constant with probability 1 on Ω . This yields the density coefficients $(\alpha(\mathbf{s}))_{\mathbf{s} \in \Sigma}$ (Lemma 8.11).
3. We show that the coefficients are positive, which completes the proof of Theorem 7.6 (Lemma 8.12).
4. Theorem 7.8 is then obtained as a corollary of Theorem 7.6.

We stress the following facts: Theorem 7.6 (LLN) does not use the properties of *distributed probabilities*: only the *homogeneity* stated in Theorem 6.4 is actually used. In contrast, the properties of distributed probabilities are used in the proof of Theorem 7.8 (extended LLN).

Prior to proceeding with the above steps of the proof, we need to introduce some new notions, that are the topic of §§8.1–8.5. Then we focus in §§8.6 on the proofs of the theorems.

8.1 An Auxiliary Markov Chain

We first associate to each Markov net a finite Markov chain that “codes”, in a sense to be made precise, the probabilistic behavior of the net. The chain is defined on a very large state space. It is thus of little interest in practice; but it has a theoretical merit, namely, in relating Markov nets to finite Markov chain theory. See a further discussion in § 8.2 below.

Denote by Q the finite set consisting of all pairs (m, z) , where $m \in \mathbf{M}_{\text{rs}}$ and $z \in \Omega_{B_0(E^m)}$. That is, z is maximal in the max-min prefix of E^m . Recall from §5.1 the definition of isomorphism of probability spaces.

Lemma 8.1 *Let $(\mathcal{N}, (p_{\mathbf{s}})_{\mathbf{s} \in \Sigma})$ be a Markov net, with Q the finite set defined as above, and let $(V_n)_{n \geq 0}$ and $(Z_n)_{n > 0}$ form the normal decomposition of maximal configurations. Then the sequence $(\gamma(V_n), Z_{n+1})_{n \geq 0}$, seen as a sequence of random variables with values in Q , defines a finite Markov chain.*

If $(\Xi, \mathfrak{G}, \mathbb{Q})$ denotes the canonical probability space associated with the Markov chain, there is an isomorphism of probability spaces $(\Omega, \mathfrak{F}, \mathbb{P}) \rightarrow (\Xi, \mathfrak{G}, \mathbb{Q})$, where \mathfrak{F} denotes the Borel σ -algebra of Ω , associating an element $\omega \in \Omega$ with the sequence $(\gamma(v_n), z_n)_{n \geq 0}$, where $(v_n, z_n)_{n \geq 0}$ denotes the normal decomposition of ω defined in §4.4.

Proof. Recall the notations γ^m, \mathbb{P}^m , etc. . . , used in reference with the net \mathcal{N}^m , for any reachable marking m . For any $m \in \mathbf{M}_{\text{rs}}$, let also Z^m be the random variable $Z^m : \Omega_{E^m} \rightarrow \Omega_{B_0(E^m)}$, $\xi \in \Omega_{E^m} \mapsto Z^m(\xi) = \xi \cap B_0(E^m)$. We define a transition matrix P on Q as follows:

$$\forall (m, z), (m', z') \in Q, \quad P_{(m,z),(m',z')} = \mathbf{I}(\gamma^m(z) = m') \mathbb{P}^{m'}(Z^{m'} = z'), \quad (31)$$

where $\mathbf{I}(\text{PRED})$ takes the value 1 if the predicate PRED is true, 0 otherwise. It is clear that P is indeed a stochastic matrix on Q , i.e.:

$$\forall (m, z) \in Q, \quad \sum_{(m', z') \in Q} P_{(m,z),(m',z')} = 1.$$

For any $n \geq 1$, let $h_n(m_0, z_1, \dots, m_{n-1}, z_n)$, with $(m_i, z_{i+1}) \in Q$ for all $i = 1, \dots, n-1$, denote the following quantity:

$$h_n(m_0, z_1, \dots, m_{n-1}, z_n) = \mathbb{P}(\gamma(V_0) = m_0, Z_1 = z_1, \dots, \gamma(V_{n-1}) = m_{n-1}, Z_n = z_n).$$

We have, by the chain rule:

$$\begin{aligned} h_n(m_0, z_1, \dots, m_{n-1}, z_n) &= h_{n-1}(m_0, z_1, \dots, m_{n-2}, z_{n-1}) \\ &\quad \times \mathbb{P}^{\gamma^{m_{n-2}}(z_{n-1})}(m(V_1) = m_{n-1}, Z^{m_{n-1}} = z_n) \\ &= h_{n-1}(m_0, z_1, \dots, m_{n-2}, z_{n-1}) \\ &\quad \times \mathbf{I}(\gamma^{m_{n-2}}(z_{n-1}) = m_{n-1}) \times \mathbb{P}^{m_{n-1}}(Z^{m_{n-1}} = z_n) \\ &= h_{n-1}(m_1, z_1, \dots, m_{n-1}, z_{n-1}) \times P_{(m_{n-2}, z_{n-1}), (m_{n-1}, z_n)}. \end{aligned}$$

This shows that the sequence $(\gamma(V_n), Z_n)_{n \geq 0}$ is a Markov chain with transition matrix P . The initial distribution μ of the chain is given by:

$$\forall (m, z) \in Q, \quad \mu(m, z) = \mathbf{I}(m = m_0) \mathbb{P}(Z_1 = z), \quad (32)$$

where m_0 is the initial marking of the net.

Let $(\Xi, \mathcal{G}, \mathbb{Q})$ be the canonical sample space associated with the Markov chain. We have by construction a measurable map $\Phi : \Omega \rightarrow \Xi$, given by $\Phi(\omega) = (\gamma(V_n(\omega)), Z_{n+1}(\omega))_{n \geq 0}$, such that $\Phi \mathbb{P} = \mathbb{Q}$, where $\Phi \mathbb{P}$ is the image probability of \mathbb{P} by Φ (see §5.1 for the definition of image probability). Let us show that Φ is injective. Indeed, if $\Phi(\omega) = (m_n, Z_{n+1})_{n \geq 0}$, then the normal decomposition $(V_n, Z_n)_{n \geq 0}$ of ω is entirely determined by $(Z_n)_{n \geq 0}$, since $V_0 = \emptyset$, and $V_n = Z_1 \oplus \dots \oplus Z_n$ for $n > 0$. Since $\omega = \sup_n V_n$ by Theorem 4.8, ω is entirely determined by $\Phi(\omega)$, which shows that Φ is injective. We finally show that Φ is onto. To this end, let $(m_n, z_{n+1})_{n \geq 0}$ be an element of Ξ . Consider the following sequence:

$$v_0 = \emptyset, \quad n > 0, \quad v_n = v_{n-1} \oplus z_n.$$

We show by induction on n that v_n is a configuration of E , with $\gamma(v_n) = m_n$. This is trivial for $n = 0$. To see it for $n = 1$, we may assume without loss of generality that $\mu(m_0, z_1) > 0$, since:

$$\mathbb{Q}(\mu(m_0, z_1) = 0) = 0.$$

According to (32), this implies in particular that $\mathbf{I}(m = m_0) = 1$, so that $z_1 \in \Omega_{B_0(E)}$. Hence $v_1 = z_1$ is a configuration of E . For the same reason, we may assume that $P_{(m_0, z_1), (m_1, z_2)} > 0$. According to (31), this implies that $\mathbf{I}(\gamma^{m_0}(z_1) = m_1) = 1$, and thus $\gamma(v_1) = m_1$, which shows the induction hypothesis for $n = 1$. The general case follows along the same line. Hence $(v_n)_{n \geq 0}$ is a nondecreasing sequence of configurations of E . Therefore $v = \sup_{n \geq 0} v_n$ is also a configuration of E . Let ω be any maximal configuration of E containing v . Then it is clear, by induction on $n \geq 0$, that the normal decomposition of ω satisfies $V_n(\omega) = v_n$ for $n \geq 0$, $Z_n(\omega) = z_n$ for $n > 0$. Hence, by Theorem 4.8, $\omega = \sup_n v_n = v$. Therefore, ω satisfies $\Phi(\omega) = (m_n, z_{n+1})_{n \geq 0}$. This shows that Φ is onto a set of \mathbb{Q} -probability 1, and completes the proof (since Φ is clearly bi-measurable). \square

8.2 Further Discussion: Markov Nets and Markov Chains

Lemma 8.1 might be surprising: after all, a Markov net can be described by means of a special Markov chain, defined on a finite—and huge—set. However, the result must not be misunderstood. We cannot just refer to Markov chain theory and invoke the lemma to end the story!

First, equivalence of probability spaces is not a strong equivalence. In particular, almost nothing on the structure of the Markov chain reveals the concurrency properties of the net, whereas they are clearly revealed by the net representation.

Second, our auxiliary Markov chain is not canonical; alternative ones could be considered. Indeed, starting from a given marking m , we could change the rule used in selecting a particular stopping prefix of E^m (here we took the max-min prefix $B_0(E^m)$). For instance we could define an arbitrary total order on dynamic clusters and pick them in this order instead of taking the max-min prefix—although we would need the integrable concurrency height for this to work.

And, last but not least, the asymptotic behavior of the net that we describe in our LLN (Theorems 7.6 and 7.8) is much more precise than what could be said about the auxiliary Markov chain. *A priori*, nothing can be said about the ergodicity of the chain. We mentioned that several other auxiliary Markov chains could have been defined; it is part of our contribution to show that, through the net representation, all of these potential Markov chains share some common asymptotic properties (this is the uniqueness of the limit involved in our LLN).

Hence, the auxiliary Markov chain really appears as a technical tool (essentially used in the proof of the LLN to get the existence of some limit). But it would not be reasonable to use it in practice to characterize the behavior of Markov nets.

Note that this auxiliary Markov chain is also different from the usual one associated with the marking graph of the net. Our auxiliary Markov chain makes the true-concurrency probabilistic semantics “rigid”, whereas the latter relates to the interleaving semantics.

Remark also that Lemma 8.1 only uses the homogeneity property (16); the fact that \mathbb{P} is a distributed product is not essential here, it is only a sufficient condition for the homogeneity. Hence the same result holds even without the local finiteness assumption, provided that the probability has the homogeneity property. In turn, the associated Markov chain would be, in general, defined on an *infinite* state space. This enlightens the role of the local finiteness assumption.

8.3 Preliminaries on Global Recurrence

Recall that we have defined in §6.3 the return operator R associated to a Markov net. Also, the successive returns are defined by:

$$R_0 = \emptyset, \quad \forall n > 0, \quad R_n(\omega) = R_{n-1}(\omega) \oplus R(\omega \ominus R_{n-1}(\omega)). \quad (33)$$

In case of a Markov net that reduces to a Markov chain, R_n coincides with the n^{th} return to the initial state. These are known to be stopping times. For general Markov nets, we have:

Lemma 8.2 *If \mathcal{N} is recurrent, the successive return operators R_n form a regular sequence of stopping operators.*

Proof. The fact that every R_n is a stopping operator follows by induction from the fact that R is a stopping operator. But this follows easily from the definitions (see a detailed proof in [9, Ch.6, l-1.6] or in [17]).

Let ω be an element of Ω such that all $R_n(\omega)$ are well-defined, $n \geq 0$. We have $\bigcup_{n \geq 0} R_n(\omega) \subseteq \omega$ by construction; for the converse inclusion, let $v = \bigcup_{n \geq 0} R_n(\omega)$, and assume that v is not maximal. Then there is an event e minimal in E^v . But then there is an integer n such that e is minimal in $E^{R_p(\omega)}$ for all $p \geq n$. In particular, e is a minimal event of both $E^{R_n(\omega)}$ and $E^{R_{n+1}(\omega)}$, which contradicts the definition of R . Hence v is maximal and thus $v = \omega$. Finally it is clear that $\langle \mathbf{1}, R \rangle \geq 1$, and thus by induction, using (33), $\langle \mathbf{1}, R_n \rangle \geq n$. This shows that $(R_n)_{n \geq 0}$ is a regular sequence of stopping operators, and completes the proof of the lemma. \square

Recall that if $\varphi : P \rightarrow Q$ is a measurable map, where P and Q are two sets respectively equipped with the σ -algebras \mathfrak{F} and \mathfrak{G} , the σ -algebra $\langle \varphi \rangle$ generated by φ is defined by $\langle \varphi \rangle = \{\varphi^{-1}(A), A \in \mathfrak{G}\}$, and then $\langle \varphi \rangle \subseteq \mathfrak{F}$.

Let \mathcal{N} be a recurrent net. We may assume without loss of generality that $R_n(\omega)$ is finite for every $n \geq 0$ (e.g., by redefining $R(\omega) = \emptyset$ whenever $R(\omega) = \omega$, note that the set of all these ω has probability 0). Hence R_n takes its values in a set at most countable. Therefore the associated σ -algebra $\langle R_n \rangle$ can be described as follows:

$$\langle R_n \rangle = \langle R_n^{-1}(u), u \in \mathcal{W}_E \rangle;$$

that is, $\langle R_n \rangle$ is the smallest σ -algebra that contains the family of subsets $\{R_n^{-1}(u), u \in \mathcal{W}_E\}$.

Furthermore, we note the following property of R_n 's, which is a general property of stopping operators: If u is a finite R -stopped configuration such that $u = R_n(\omega_o)$ for some $n \geq 0$ and $\omega_o \in \Omega$, then we have:

$$R_n^{-1}(u) = \{\omega \in \Omega : \omega \supseteq u\}. \quad (34)$$

Indeed, since R_n is a stopping operator by Lemma 8.2, it is enough to verify (34) for stopping operators; but this is an immediate consequence of point 2 in Definition 7.2.

Corollary 8.3 *Assume that \mathcal{N} is recurrent, and denote by \mathfrak{F} the Borel σ -algebra on Ω , and for all $n \geq 0$, by \mathfrak{F}_n the σ -algebra generated by R_n . Then $\mathfrak{F} = \langle \mathfrak{F}_n, n \geq 0 \rangle$, i.e., \mathfrak{F} coincides with the smallest σ -algebra that contains all $\mathfrak{F}_n, n \geq 0$.*

Proof. Let $\mathfrak{G} = \langle \mathfrak{F}_n, n \geq 0 \rangle$. We obviously have $\mathfrak{G} \subseteq \mathfrak{F}$. For the converse inclusion, recall the notation $\mathcal{S}(u) = \{\omega \in \Omega : \omega \supseteq u\}$. Since \mathfrak{F} is generated by the collection $\{\mathcal{S}(u)\}$, where u ranges over the set \mathcal{V}_E of finite configurations of E , it is enough to show that $\mathcal{S}(u) \in \mathfrak{G}$ for every $u \in \mathcal{V}_E$.

Let $K = \{R_n(\omega), n \geq 0, \omega \in \Omega\}$. K is at most countable, since it consists of finite configurations. Let $u \in \mathcal{V}_E$. From the equality $\omega = \sup_{n \geq 0} R_n(\omega)$ stated in

Lemma 8.2, we get:

$$\mathcal{S}(u) = \bigcup_{\substack{v \in K \\ v \supseteq u}} \mathcal{S}(v).$$

Using (34), each $\mathcal{S}(v)$ with $v \in K$ can be written $\mathcal{S}(v) = R_{n_v}^{-1}(v)$, where $n_v \geq 0$. Thus:

$$\mathcal{S}(u) = \bigcup_{\substack{v \in K \\ v \supseteq u}} R_{n_v}^{-1}(v),$$

a union at most countable of \mathfrak{G} -measurable subsets. Hence $\mathcal{S}(u) \in \mathfrak{G}$, which completes the proof of the corollary. \square

8.4 Preliminaries on Local Recurrence

Fix \mathbf{s}_0 a dynamic cluster of \mathcal{N} , and consider $\omega \in \Omega$. Assume that x and x' are two branching cells in $\Delta(\omega)$ such that $\langle x \rangle = \langle x' \rangle = \mathbf{s}_0$. Then there are events in x and x' that are both compatible, and labeled by the same transition. Since the net is safe, this implies that these two events are causally related. This induces in turn an ordering of such branching cells, for \mathbf{s}_0 and ω fixed, which is thus a total ordering. Furthermore, for each branching cell $x \in \Delta(\omega)$, the configuration:

$$v_x = \inf\{v \in \mathcal{W}, v \subseteq \omega, x \in \delta_E(v)\}, \quad (35)$$

is a finite R -stopped configuration, since this set of configurations is nonempty, and by the lattice property of compatible R -stopped sub-configurations of ω (Theorem 4.7, point 5). The above ordering on branching cells x such that $x \in \Delta(\omega)$ and $\langle x \rangle = \mathbf{s}_0$ corresponds to the set-inclusion of associated configurations v_x .

In particular, again with ω fixed, if the following set is nonempty:

$$\{x \in \Delta(\omega) : \langle x \rangle = \mathbf{s}_0, x \notin \delta_E(\emptyset)\},$$

it has a unique minimal element x , with an associated configuration v_x defined as in (35). Let $S^{\mathbf{s}_0}(\omega)$ denote this configuration, so that $S^{\mathbf{s}_0}(\omega)$ is defined by:

$$S^{\mathbf{s}_0}(\omega) = \inf\{v \in \mathcal{W} : v \subseteq \omega, v \neq \emptyset, \exists x \in \delta_E(v), \langle x \rangle = \mathbf{s}_0\}. \quad (36)$$

Remark that the branching cell x in (36) is then unique, again for safeness reasons. Moreover, since the net is assumed recurrent, a simple Borel-Cantelli argument shows that $S^{\mathbf{s}_0}$ is defined on Ω with probability 1.

Definition 8.4 (local return operator) *If the Markov net $(\mathcal{N}, (p_s)_{s \in \Sigma})$ is recurrent, for every dynamic cluster \mathbf{s}_0 , the local return to \mathbf{s}_0 is the mapping $S^{\mathbf{s}_0} : \Omega \rightarrow \mathcal{W}$ defined by (36), with probability 1.*

Intuitively, local returns are returns with “minimal moves of the tokens”. For every reachable marking m , and in particular if m is recursively stopped, the same definition of local return operator applies to Ω_{E^m} . Since \mathcal{N}^m is itself also recurrent, the local

return to s_0 is also defined on Ω_{E^m} with \mathbb{P}^m -probability 1. We denote this local return operator by:

$$S^{\mathbf{s}_0, m} : \Omega_{E^m} \rightarrow \mathcal{W}_{E^m}, \quad \text{with } \mathbb{P}^m\text{-probability 1.} \quad (37)$$

This allows to construct the *successive local return operators* to s_0 as follows. The local returns $(S_n^{\mathbf{s}_0})_{n \geq 1}$ to s_0 are defined on Ω with probability 1 by:

$$S_1^{\mathbf{s}_0} = S^{\mathbf{s}_0}, \quad \forall n > 1, \quad S_n^{\mathbf{s}_0}(\omega) = S_{n-1}^{\mathbf{s}_0}(\omega) \oplus S^{\mathbf{s}_0, m_{n-1}}(\omega \ominus S_{n-1}^{\mathbf{s}_0}(\omega)), \quad (38)$$

with $m_{n-1} = \gamma(S_{n-1}^{\mathbf{s}_0}(\omega))$.

We summarize the properties of local return operators as follows:

Proposition 8.5 *Let $(\mathcal{N}, (p_s)_{s \in \Sigma})$ be a recurrent Markov net, and let \mathbf{s}_0 be a dynamic cluster of \mathcal{N} . Then for each $n \geq 1$, the n^{th} local return operator $S_n^{\mathbf{s}_0}$ to \mathbf{s}_0 is a stopping operator of Ω . If $(\mathcal{N}, (p_s)_{s \in \Sigma})$ has integrable concurrency height, then $(S_n^{\mathbf{s}_0})_{n \geq 1}$ is a regular sequence of stopping operators.*

Proof. As for global return operators, it is enough to show that the single $S^{\mathbf{s}_0}$ defined by (36) is a stopping operator. To this end, we clearly have $S^{\mathbf{s}_0}(\omega) \subseteq \omega$. Let $\omega, \omega' \in \Omega$, set $v = S^{\mathbf{s}_0}(\omega)$, $v' = S^{\mathbf{s}_0}(\omega')$, and assume that $\omega' \supseteq v$. Let x be the unique element of $\delta_E(v)$ such that $\langle x \rangle = \mathbf{s}_0$. Then, since v is finite nonempty R -stopped, since $x \in \delta_E(v)$ and since $v \subseteq \omega'$, we have $v' \subseteq v$ by the very definition of $v' = S^{\mathbf{s}_0}(\omega')$. This implies $v' \subseteq \omega$, and in turn, by minimality of $S^{\mathbf{s}_0}(\omega)$, we get $v \subseteq v'$. Finally, $v = v'$, which proves that $S^{\mathbf{s}_0}$ is a stopping operator.

Assume moreover that the net has integrable concurrency height; we show that the sequence $(S_n^{\mathbf{s}_0})_{n \geq 1}$ of local returns to \mathbf{s}_0 is regular. $(S_n^{\mathbf{s}_0})_{n \geq 1}$ is clearly nondecreasing. Since $S^{\mathbf{s}_0} \neq \emptyset$, $\langle \mathbf{1}, S(\omega) \rangle \geq 1$ with probability 1; it follows thus from (38) that $\langle \mathbf{1}, S_n^{\mathbf{s}_0} \rangle \geq n$ for all $n \geq 1$. Hence it remains only to show that $\sup_{n \geq 1} S_n^{\mathbf{s}_0}(\omega) = \omega$ with probability 1. Fix $\omega \in \Omega$, let $v_n = S_n^{\mathbf{s}_0}(\omega)$, $v = \sup_{n \geq 1} v_n$, and assume that $v \neq \omega$. Consider any minimal event $e \in E^v$. Then there is an integer p such that e is also a minimal event of E^{v_n} for all $n \geq p$. Pick any place b in the preset of $\lambda(e)$, where $\lambda(e)$ is the transition that labels e . For $m = \gamma(v_p)$, the token in the place b is a frozen token, from v_p to ω , i.e., the length $\langle \mathbf{1}, \omega \ominus v_p \rangle$ is infinite. But, since the length is integrable, all such ω have together probability 0. This shows that $\sup_{n \geq 1} S_n^{\mathbf{s}_0}(\omega) = \omega$ with probability 1, and completes the proof of the proposition. \square

Still consider a fixed dynamic cluster \mathbf{s}_0 , and let $\omega \in \Omega$ be such that the local returns $S_n^{\mathbf{s}_0}(\omega)$ to \mathbf{s}_0 are defined for all $n \geq 1$, which holds with probability 1 if the net is recurrent. For each $n \geq 1$, there is by construction a unique branching cell $x_n(\omega) \in \delta_E(S_n^{\mathbf{s}_0}(\omega))$ such that $\langle x_n(\omega) \rangle = \mathbf{s}_0$. As a consequence of Lemma 4.3, we have $\omega \cap x_n(\omega) \in \Omega_x = \Omega_{\mathbf{s}_0}$. Therefore, if we set:

$$\forall n \geq 1, \quad Y_n^{\mathbf{s}_0}(\omega) = \omega \cap x_n(\omega), \quad (39)$$

we get a sequence $(Y_n^{\mathbf{s}_0})_{n \geq 1}$ of random variables such that:

$$\forall n \geq 1, \quad Y_n^{\mathbf{s}_0}(\omega) \in \Omega_{\mathbf{s}_0}, \quad \text{with probability 1.}$$

Using 1) the Strong Markov property for Markov nets, and 2) the properties of the distributed product, yields the following result regarding the sequence $(Y_n^{s_0})_{n \geq 1}$ (see [9, Ch.6, III-3.2] for the proof):

Lemma 8.6 *Let $(\mathcal{N}, (p_s)_{s \in \Sigma})$ be a recurrent Markov net, let s_0 be a dynamic cluster, and let $(Y_n^{s_0})_{n \geq 1}$ be the sequence of random variables, with values in Ω_{s_0} , defined by (39). Then $(Y_n^{s_0})_{n \geq 1}$ is a sequence of independent and identically distributed variables, with law p_{s_0} in Ω_{s_0} .*

We finally show the following result, which generalizes to Markov nets the *positive recurrence* of recurrent states, for Markov chains:

Lemma 8.7 (positive recurrence) *Let $(\mathcal{N}, (p_s)_{s \in \Sigma})$ be a recurrent Markov net. Then the global return R , and all the local returns $S^{s,m}$, $s \in \Sigma$, $m \in \mathbf{M}_{rs}$, satisfy:*

$$\mathbb{E}(\langle \mathbf{1}, R \rangle) < \infty, \quad \mathbb{E}^m(\langle \mathbf{1}, S^{s,m} \rangle) < \infty,$$

where \mathbb{E} and \mathbb{E}^m denote respectively the mathematical expectation w.r.t. probabilities \mathbb{P} and \mathbb{P}^m .

Proof. We first show that $\mathbb{E}(\langle \mathbf{1}, R \rangle) < \infty$. Let T be the integer random variable $T = \langle \mathbf{1}, R \rangle$. Recall the usual identity $\mathbb{E}(T) = \sum_{n > 0} \mathbb{P}(T \geq n)$. Assume that there is an integer $r > 0$ such that:

$$\sum_{n \geq 0} \mathbb{P}(T \geq nr) < \infty. \tag{40}$$

Then we have:

$$\begin{aligned} \mathbb{E}(T) &= \sum_{n > 0} \mathbb{P}(T \geq n) = \sum_{i=0}^{\infty} \sum_{j=1}^r \mathbb{P}(T \geq ir + j) \\ &\leq \sum_{i=0}^{\infty} \sum_{j=1}^r \mathbb{P}(T \geq ir) \\ &= r \sum_{n=0}^{\infty} \mathbb{P}(T \geq nr) < \infty, \quad \text{by (40)}. \end{aligned}$$

Hence, to prove $\mathbb{E}(T) < \infty$, it is enough to show (40). Let $(V_n)_{n \geq 0}$ be the sequence of stopping operators coming from the normal decomposition of maximal configurations, as defined in §4.4 and in §8.1, Lemma 8.1. There is an integer $k_1 > 0$ such that $\langle \mathbf{1}, V_n \rangle \leq k_1 \times n$ for all $n > 0$; take for example k_1 as the maximal number of simultaneously enabled transitions of the net. Define the random variable T' by:

$$T' = \inf\{n \geq 0 : R \subseteq V_n\}.$$

Since R is finite with probability 1, and since $\bigcup_n V_n(\omega) = \omega$ for all $\omega \in \Omega$, T' is finite with probability 1. Therefore, with probability 1, $T \leq \langle \mathbf{1}, V_{T'} \rangle \leq k_1 \times T'$. Hence, to show (40), it is enough to show that there is an integer $q > 0$ such that:

$$\sum_{n \geq 0} \mathbb{P}(T' \geq qn) < \infty. \tag{41}$$

We choose the integer q as follows. Recall that \mathbf{M}_{rs} denotes the set of recursively stopped markings of the net. Since the net is recurrent, we choose for any marking $m \in \mathbf{M}_{\text{rs}}$ a finite R -stopped configuration v_m leading back from m to the initial marking, after having moved all tokens in the net. Each v_m has \mathbb{P}^m -positive likelihood. We put $q = \max_{m \in \mathbf{M}_{\text{rs}}} \langle \mathbf{1}, v_m \rangle$, which is finite since \mathbf{M}_{rs} is a finite set, and positive otherwise we would have $v_m = \emptyset$ for all $m \in \mathbf{M}_{\text{rs}}$. Let n be any integer $n \geq 1$, and let \mathbb{Q} be the conditional probability $\mathbb{Q} = \mathbb{P}(\cdot | T' \geq q(n-1))$. Then we have, by the Bayes rule:

$$\mathbb{Q}(T' \geq qn) = \sum_{m \in \mathbf{M}_{\text{rs}}} \mathbb{Q}(T' \geq qn | \gamma(V_{q(n-1)}) = m) \times \mathbb{Q}(\gamma(V_{q(n-1)}) = m), \quad (42)$$

where $\gamma(v)$ denotes as usual the marking reached by a configuration v . Denote, for any $m \in \mathbf{M}_{\text{rs}}$, by $(V_j^m)_{j \geq 0}$ the sequence of stopping operators coming from the normal decomposition of maximal configurations, defined on Ω_{E^m} . Then, using the homogeneity property (Theorem 6.4), $\mathbb{Q}(T' \geq qn | \gamma(V_{q(n-1)}) = m)$ is the \mathbb{P}^m probability that V_q^m does not contain any return to the initial marking. By definition of q , this probability is less than a constant $a < 1$. Hence we get from (42) that $\mathbb{Q}(T' \geq qn) \leq a$, and thus, coming back to the definition of \mathbb{Q} :

$$\mathbb{P}(T' \geq qn) \leq a \mathbb{P}(T' \geq q(n-1)) \leq a^2 \mathbb{P}(T' \geq q(n-2)) \leq \dots \leq a^n.$$

Since $a < 1$, Equation (41) follows, hence $\mathbb{E}(T) < \infty$.

We now show that $\mathbb{E}(\langle \mathbf{1}, S^{\mathbf{s}, m} \rangle) < \infty$ for every $\mathbf{s} \in \Sigma$ and $m \in \mathbf{M}_{\text{rs}}$. Since the net \mathcal{N}^m is recurrent by Lemma 6.8, we may assume without loss of generality that $m = m_0$ is the initial marking. If there is a branching cell x such that: x is an initial stopping prefix of E , and such that $\langle x \rangle = \mathbf{s}$, then $S^{\mathbf{s}, m_0} \subseteq R$. Therefore $\langle \mathbf{1}, S^{\mathbf{s}, m_0} \rangle \leq \langle \mathbf{1}, R \rangle$ and thus $\mathbb{E}(\langle \mathbf{1}, S^{\mathbf{s}, m_0} \rangle) \leq \mathbb{E}(\langle \mathbf{1}, R \rangle) < \infty$.

Finally we show that the general case reduces to this particular case. For any $M \in \mathbf{M}_{\text{rs}}$, define the M -reachability operator A^M by:

$$\forall \xi \in \Omega_{E^m}, \quad A^M(\omega) = \min\{v \in \mathcal{W}_E : v \subseteq \xi, \gamma(v) = M, \text{Min}(E) \cap \text{Min}(E^v) = \emptyset\}.$$

Then A^M is finite with probability 1, and using the very same technique than above, we conclude that $\mathbb{E}(\langle \mathbf{1}, A^M \rangle) < \infty$. Now for any $\mathbf{s} \in \Sigma$, let $M \in \mathbf{M}_{\text{rs}}$ such that there is a branching cell x and a $v \in \mathcal{W}_E$ with $\gamma(v) = M$, $\langle x \rangle = \mathbf{s}$ and $x \in \delta_E(v)$. Then we have:

$$\forall \omega \in \Omega, \quad S^{\mathbf{s}, m_0}(\omega) \subseteq A^M(\omega) \oplus S^{M, \mathbf{s}}(\omega \ominus A^M(\omega)).$$

Therefore:

$$\mathbb{E}(\langle \mathbf{1}, S^{\mathbf{s}, m_0} \rangle) \leq \mathbb{E}(\langle \mathbf{1}, A^M \rangle) + \mathbb{E}^M(\langle \mathbf{1}, S^{\mathbf{s}, M} \rangle).$$

We have $\mathbb{E}(\langle \mathbf{1}, A^M \rangle) < \infty$ as we remarked above, and $\mathbb{E}^M(\langle \mathbf{1}, S^{\mathbf{s}, M} \rangle) < \infty$ by the first case. Therefore $\mathbb{E}(\langle \mathbf{1}, S^{\mathbf{s}, m_0} \rangle) < \infty$, which completes the proof of the lemma. \square

8.5 Preliminaries on Ergodicity

Definition 8.8 (homogeneous function) Let \mathbf{M}_{rs} denote as usual the set of recursively stopped markings of \mathcal{N} —i.e., those markings reached by some finite R -stopped configurations. A family $H = (H^m)_{m \in \mathbf{M}_{rs}}$ of real-valued functions $H^m : \Omega_{E^m} \rightarrow \mathbb{R} \cup \{\infty\}$ is called a homogeneous function.

The homogeneous function H is said to be nonnegative, respectively integrable, if every H^m is nonnegative, respectively integrable w.r.t. probability \mathbb{P}^m .

If $H = (H^m)_{m \in \mathbf{M}_{rs}}$ is a homogeneous function, we consider the family $(H^v)_{v \in \mathcal{W}_E}$ of real-valued functions $H^v : \Omega_{E^v} \rightarrow \mathbb{R} \cup \{\infty\}$ defined by:

$$\forall v \in \mathcal{W}_E, \quad H^v = H^{\gamma(v)},$$

and that satisfies:

$$\forall v, v' \in \mathcal{W}, \quad \gamma(v) = \gamma(v') \Rightarrow H^v = H^{v'}.$$

The two representations $(H^m)_{m \in \mathbf{M}_{rs}}$ and $(H^v)_{v \in \mathcal{W}}$ are obviously equivalent.

We now prove the following result, to be interpreted as an ergodicity result.

Theorem 8.9 Let \mathcal{N} be a recurrent Markov net. Let $H = (H^v)_{v \in \mathcal{W}}$ be a nonnegative and integrable homogeneous function. Assume that we have, with probability 1:

$$\forall \omega \in \Omega, \quad \forall v \in \mathcal{W}_E, \quad v \subseteq \omega \Rightarrow H^v(\omega \ominus v) = H^\emptyset(\omega). \quad (43)$$

Then for every $v \in \mathcal{W}_E$, H^v is constant on a set of probability 1, and all the so obtained constants are identical.

Proof. We begin with the following observation: If V is a stopping operator, with $V(\omega)$ finite with probability 1, and if $\mathbb{E}(\cdot | V)$ denotes the conditional expectation w.r.t. the σ -algebra $\langle V \rangle$ generated by V (see §5.1), we have for every integrable function $f : \Omega \rightarrow \mathbb{R}$:

$$\mathbb{E}(f | V) = \int_{\Omega^{\gamma(V)}} f(V \oplus \xi) d\mathbb{P}^{\gamma(V)}(\xi), \quad (44)$$

where $V \oplus \xi$ denotes the concatenation of V and ξ (proof left to the reader; hint: use (34) with stopping operator V).

We show first that H^\emptyset is constant on Ω with probability 1. By Corollary 8.3, we have $\mathfrak{F} = \langle \mathfrak{F}_n, n \geq 0 \rangle$, where $\mathfrak{F}_n = \langle R_n \rangle$ and \mathfrak{F} is the Borel σ -algebra on Ω . The Martingale convergence theorem [10, Th. 35.5] implies thus:

$$H^\emptyset = \lim_{n \rightarrow \infty} \mathbb{E}(H^\emptyset | \mathfrak{F}_n), \quad \text{with probability 1.} \quad (45)$$

We apply (44) with $f = H^\emptyset$ to get:

$$\begin{aligned} \mathbb{E}(H^\emptyset | \mathfrak{F}_n) &= \int_{\Omega^{\gamma(R_n)}} H^\emptyset(R_n \oplus \xi) d\mathbb{P}^{\gamma(R_n)}(\xi) \\ &= \int_{\Omega} H^\emptyset(R_n \oplus \xi) d\mathbb{P}(\xi), \end{aligned} \quad (46)$$

the latter by applying the homogeneity of \mathbb{P} (Theorem 6.4), and since $\gamma(R_n) = m_0$ by construction of the successive return operators R_n . Applying the assumption (43) with $v = R_n$, we get:

$$\forall \xi \in \Omega, \quad H^\theta(R_n \oplus \xi) = H^{R_n}(\xi) = H^\theta(\xi), \quad (47)$$

the latter since H is a homogeneous function. Using together (45), (46) and (47), we obtain:

$$H^\theta(\omega) = \lim_{n \rightarrow \infty} \int_{\Omega} H^\theta(R_n \oplus \xi) d\mathbb{P}(\xi) = \int_{\Omega} H^\theta(\xi) d\mathbb{P}(\xi) = \mathbb{E}(H^\theta),$$

with probability 1. This shows that H^θ is constant on Ω with probability 1. For the same reasons, every H^v , with $v \in \mathcal{W}$, is constant on Ω^v with \mathbb{P}^v probability 1. It follows from (43), and from the fact that the likelihood of every $v \in \mathcal{W}$ is positive, that all the constants coincide. \square

Finally, the following result on homogeneous functions will be useful. It is the adaptation of a classical lemma from dynamical systems theory; our proof is borrowed from [18].

Proposition 8.10 *Let \mathcal{N} be a Markov net with unfolding E , and let $(V_n)_{n \geq 0}$ be a regular sequence of stopping operators on Ω . For each $n \geq 0$, we define θ_n by:*

$$\forall \omega \in \Omega, \quad \theta_n(\omega) = \omega \ominus V_n(\omega), \quad \text{so that : } \theta_n \in \Omega_{E^{\gamma(V_n)}}.$$

Let also $H = (H^v)_{v \in \mathcal{W}}$ be a homogeneous, nonnegative and integrable function. Then we have:

$$\lim_{n \rightarrow \infty} \frac{H^{V_n}(\theta_n)}{\langle \mathbf{1}, V_n \rangle} = 0, \quad \text{with probability 1.}$$

Proof. We also use the notation $(H^m)_{m \in \mathbf{M}_{rs}}$ for H , where \mathbf{M}_{rs} is the set of recursively stopped markings. Since $(V_n)_{n \geq 0}$ is a regular sequence of operators, there is a constant $k > 0$ such that $\langle \mathbf{1}, V_n \rangle \geq k \times n$ for every $n \geq 1$. Therefore it is enough to show:

$$\lim_{n \rightarrow \infty} \frac{H^{V_n}(\theta_n)}{n} = 0, \quad \text{with probability 1.}$$

To this end, we denote by X_n the random variable $X_n = \frac{1}{n} H^{V_n}(\theta_n)$, and we use the following classical criterion that implies the convergence of $(X_n)_{n \geq 1}$ to 0 with probability 1:

$$\forall \epsilon > 0, \quad \sum_{n \geq 1} \mathbb{P}(X_n \geq \epsilon) < \infty. \quad (48)$$

For each $n \geq 1$, let K_n denote the set of values of V_n . Since V_n is finite with probability 1, we assume without loss of generality that K_n is at most countable. Since V_n is a stopping operator, we have the property already observed:

$$\forall u \in K_n, \quad \{\omega \in \Omega : V_n(\omega) = u\} = \mathcal{S}(u),$$

where $\mathcal{S}(u)$ denotes as usual $\mathcal{S}(u) = \{\omega \in \Omega : \omega \supseteq u\}$. From this we get:

$$\begin{aligned}
 \mathbb{P}(X_n \geq \epsilon) &= \sum_{u \in K_n} \mathbb{P}(V_n = u) \mathbb{P}(H^{V_n}(\theta_n) \geq n\epsilon \mid V_n = u) \\
 &= \sum_{u \in K_n} \mathbb{P}(V_n = u) \mathbb{P}^u(H^u \geq n\epsilon) \\
 &= \sum_{m \in \mathbf{M}_{rs}} \left(\sum_{\substack{u \in K_n \\ \gamma(u)=m}} \mathbb{P}(V_n = u) \right) \mathbb{P}^m(H^m \geq n\epsilon) \\
 &\leq \sum_{m \in \mathbf{M}_{rs}} \mathbb{P}^m(H^m \geq n\epsilon). \tag{49}
 \end{aligned}$$

In order to show (48), and from (49), it is enough to show that for each $m \in \mathbf{M}_{rs}$, the following sum is finite: $\sum_{n \geq 1} \mathbb{P}^m(H^m \geq n\epsilon) < \infty$. Recall the usual equality $\mathbb{E}(f) = \sum_{k=1}^{\infty} \mathbb{P}(f \geq k)$ for every nonnegative integrable function $f : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$. With this transformation, we get for every $m \in \mathbf{M}_{rs}$:

$$\begin{aligned}
 \sum_{n \geq 1} \mathbb{P}^m(H^m \geq n\epsilon) &= \sum_{n \geq 1} \mathbb{P}^m\left(\frac{1}{\epsilon} H^m \geq n\right) \\
 &= \frac{1}{\epsilon} \mathbb{E}^m(H^m) < \infty,
 \end{aligned}$$

since H^m is integrable for every $m \in \mathbf{M}_{rs}$. This completes the proof of the proposition. \square

8.6 Proof of Theorems 7.6 and 7.8.

We begin with the proof of Theorem 7.6. We decompose the proof in two steps: first, the existence of the density coefficients (points 1 and 2 of the theorem, shown in Lemma 8.11 below), then their properties (point 3 of the theorem, shown in Lemma 8.12 below).

Existence of the Density Coefficients

The ergodic means $M(f, v)$ are linear in f . State functions form a vector space of finite dimension, with basis the collection of state function $\mathbf{1}_{s_0}$, $s_0 \in \Sigma$, defined by:

$$\forall s \in \Sigma, \quad \mathbf{1}_{s_0}(s) = \begin{cases} 1, & \text{if } s = s_0, \\ 0, & \text{otherwise.} \end{cases} \tag{50}$$

Therefore, to obtain the convergence of ergodic means of any state function f stated in (24) in Theorem 7.6, it is enough to show the following lemma:

Lemma 8.11 *For every dynamic cluster s_0 , the ergodic means $M(\mathbf{1}_{s_0}, \cdot)$ converge in the sense of Definition 7.3 to a function $\alpha(s_0, \cdot) : \Omega \rightarrow \mathbb{R}$. The function $\alpha(s_0, \cdot)$ is constant with probability 1 on Ω .*

Proof. We fix a dynamic cluster s_0 of \mathcal{N} , and we consider the state function $f_0 = \mathbf{1}_{s_0}$ defined by (50).

Let $(M_n, Z_{n+1})_{n \geq 0}$ denote the auxiliary Markov chain of the net (see §8.1), where $M_n = \gamma(V_n)$ is the marking associated with the configuration V_n , arising from the normal decomposition of maximal configurations. The ergodic means relative to f_0 and V_n satisfy:

$$M(f_0, V_n) = \frac{\langle f_0, V_n \rangle}{n} \frac{n}{\langle \mathbf{1}, V_n \rangle} = \underbrace{\frac{\sum_{k=0}^{n-1} \mathbf{I}(s_0 \in \delta_{E^{M_k}}(\emptyset))}{n}}_{\text{ergodic means for } (M_n, Z_n)} \cdot \left(\underbrace{\frac{\sum_{k=0}^{n-1} \text{Card}(\delta_{E^{M_k}}(\emptyset))}{n}}_{\text{ergodic means for } (M_n, Z_n)} \right)^{-1},$$

where we recall that $\mathbf{I}(\text{PRED})$ takes the value 1 if the predicate PRED is true, 0 otherwise. Hence, each factor is given by ergodic means relative to the auxiliary Markov chain $(M_n, Z_{n+1})_{n \geq 0}$. The ergodic theory of Markov chains implies that each of these factors has a limit with probability 1, according to Theorem A.1 in Appendix A. Therefore, the following real-valued random variable $G : \Omega \rightarrow \mathbb{R}$ is well defined with probability 1:

$$G(\omega) = \lim_{n \rightarrow \infty} M(f_0, V_n(\omega)).$$

For each $v \in \mathcal{W}$, the same construction applies to the probabilistic future (E^v, \mathbb{P}^v) . This defines a collection of measurable maps

$$H^v : \Omega^v \rightarrow \mathbb{R}, \tag{51}$$

with $H^\emptyset = G$. By construction, the family $H = (H^v)_{v \in \mathcal{W}}$ is a homogeneous function (Definition 8.8). According to point 2 of Lemma 8.13 below, H satisfies:

$$\forall v \in \mathcal{W}, \quad v \subseteq \omega \Rightarrow H^v(\omega \ominus v) = H^\emptyset(\omega), \quad \text{with probability 1.}$$

Combined with Theorem 8.9, this implies that H^\emptyset is constant on Ω . Let $\alpha(s_0)$ denote this constant. Then, according to point 1 of Lemma 8.13 below, for every sequence $(W_n)_{n \geq 0}$ of stopping operators, we have:

$$\lim_{n \rightarrow \infty} M(f_0, W_n) = H^\emptyset = \alpha(s_0), \quad \text{with probability 1.}$$

This shows that the ergodic means $M(f, \cdot)$ converge to the constant $\alpha(s_0)$. □

With Lemma 8.11, we have shown points 1 and 2 of Theorem 7.6. It remains to show point 3, which is the topic of next lemma.

Lemma 8.12 *The density coefficients $\alpha(s)$ satisfy:*

$$\forall s \in \Sigma, \quad 0 < \alpha(s) \leq 1, \quad \text{and} \quad \sum_{s \in \Sigma} \alpha(s) = 1. \tag{52}$$

Proof. We first show that the $\alpha(\mathbf{s})$'s sum to 1. Consider any regular sequence of stopping operators $(V_n)_{n \geq 0}$ —we have seen the existence of such a sequence in Lemma 7.4. We have obviously:

$$\forall n \geq 0, \quad \sum_{\mathbf{s} \in \Sigma} \langle \mathbf{1}_{\mathbf{s}}, V_n \rangle = \langle \mathbf{1}, V_n \rangle,$$

where $\mathbf{1}$ is the extended state function that counts all branching cells. Therefore, taking the ratio and then the limit, we get:

$$1 = \lim_{n \rightarrow \infty} \frac{1}{\langle \mathbf{1}, V_n \rangle} \sum_{\mathbf{s} \in \Sigma} \langle \mathbf{1}_{\mathbf{s}}, V_n \rangle = \lim_{n \rightarrow \infty} \sum_{\mathbf{s} \in \Sigma} M(\mathbf{1}_{\mathbf{s}}, V_n) = \sum_{\mathbf{s} \in \Sigma} \alpha(\mathbf{s}).$$

This shows that the $\alpha(\mathbf{s})$'s sum to 1, as claimed. This also implies that $0 \leq \alpha(\mathbf{s}) \leq 1$ for every $\mathbf{s} \in \Sigma$.

It remains only to show that $\alpha(\mathbf{s}) > 0$ for every $\mathbf{s} \in \Sigma$. Let $\mathbf{s} \in \Sigma$, and consider the sequence $(S_n^{\mathbf{s}})_{n \geq 1}$ of local returns to \mathbf{s} . This sequence is a regular sequence of stopping operators according to Proposition 8.5, §8.4. We have thus, applying the LLN to the extended state function $\mathbf{1}_{\mathbf{s}}$:

$$\alpha(\mathbf{s}) = \lim_{n \rightarrow \infty} \frac{1}{\langle \mathbf{1}, S_n^{\mathbf{s}} \rangle} \langle \mathbf{1}_{\mathbf{s}}, S_n^{\mathbf{s}} \rangle = \lim_{n \rightarrow \infty} \frac{n}{\langle \mathbf{1}, S_n^{\mathbf{s}} \rangle}. \quad (53)$$

The latter ratio can be written as follows:

$$\frac{n}{\langle \mathbf{1}, S_n^{\mathbf{s}}(\omega) \rangle} = \frac{n}{\langle \mathbf{1}, S_1^{\mathbf{s}}(\omega) \rangle + \sum_{k=1}^{n-1} (\langle \mathbf{1}, S_{k+1}^{\mathbf{s}}(\omega) \rangle - \langle \mathbf{1}, S_k^{\mathbf{s}}(\omega) \rangle)}.$$

We have for every $k \geq 1$:

$$\langle \mathbf{1}, S_{k+1}^{\mathbf{s}} \rangle - \langle \mathbf{1}, S_k^{\mathbf{s}} \rangle = \langle \mathbf{1}, S_{k+1}^{\mathbf{s}} \ominus S_k^{\mathbf{s}} \rangle = \langle \mathbf{1}, S^{\mathbf{s}, m_k} \rangle, \quad (54)$$

where m_k denotes the marking $\gamma(S_k^{\mathbf{s}})$, and $S^{\mathbf{s}, m}$ denotes as in (37) the local return to \mathbf{s} defined on Ω_{E^m} . It follows from the Markov property for Markov nets [17] that the sequence $(S^{\mathbf{s}, m_k})_{k \geq 1}$ is a sequence of independent random variables. They are not identically distributed since the law of $S^{\mathbf{s}, m_k}$ depends on $m_k = \gamma(S_k^{\mathbf{s}})$. But these laws range over a finite set, since markings are finitely many. Furthermore, each expectation $\mathbb{E}(\langle \mathbf{1}, S^{\mathbf{s}, m} \rangle)$ is finite according to Lemma 8.7. Therefore, we shall apply Proposition A.3 of Appendix A to get, using (53) and (54):

$$\frac{1}{\alpha(\mathbf{s})} \leq \max_m \mathbb{E}^m \langle \mathbf{1}, S^{\mathbf{s}, m} \rangle < \infty.$$

This shows that $\alpha(\mathbf{s}) > 0$. □

With Lemmas 8.11 and 8.12, we have completed the proof of Theorem 7.6. In the proof of Lemma 8.11, we have used the following lemma:

Lemma 8.13 *Assume that \mathcal{N} is a recurrent Markov net with integrable concurrency height, and let H be the homogeneous function defined in (51). Consider the state function $f_0 = \mathbf{1}_{\mathbf{s}_0}$ defined by (50). Then H satisfies:*

1. For every regular sequence $(W_n)_{n \geq 0}$ of stopping operators, we have:

$$\lim_{n \rightarrow \infty} M(f_0, W_n) = H^0, \quad \text{with probability 1.}$$

2. For each $v \in \mathcal{W}$, and with probability 1:

$$v \subseteq \omega \Rightarrow H^v(\omega \ominus v) = H^0(\omega).$$

Proof. Let $(W_n)_{n \geq 0}$ be a regular sequence of stopping operators. $(V_j, Z_{j+1})_{j \geq 0}$ denotes as above the normal decomposition of maximal configurations. For each $n \geq 1$, let $J_n : \Omega \rightarrow \mathbb{N}$ be the integer-valued random variable defined by:

$$J_n(\omega) = \inf\{p \geq 0 : V_p(\omega) \supseteq W_n(\omega)\} < \infty, \quad \text{with probability 1.}$$

Then, with probability 1, $\delta_E(V_{J_n}(\omega)) \cap \delta_E(W_n(\omega)) \neq \emptyset$. Therefore $E^{W_n(\omega)}$ and $E^{V_{J_n}(\omega)}$ possess a minimal event in common. Thus $V_{J_n}(\omega)$ is a sub-configuration of ω that keeps a token frozen, starting from the sub-configuration $W_n(\omega)$. Therefore, by definition of the integer random variables $L_m(b, \cdot)$, we have:

$$\langle \mathbf{1}, V_{J_n} \ominus W_n \rangle \leq \Gamma_{\gamma(W_n)}, \quad (55)$$

where $\Gamma_m(\xi) = \sup_{b \in m} L_m(b, \xi)$ for all recursively stopped markings m and $\xi \in \Omega_{E^m}$. Remark that, since all $L_m(b, \cdot)$ are integrable, and since there are only finitely many of them, $\Gamma = (\Gamma_m)_{m \in \mathcal{M}_{rs}}$ is a homogeneous integrable function.

We now show that the quantity:

$$\epsilon_n = M(f_0, W_n) - \frac{\langle f_0, V_{J_n} \rangle}{\langle \mathbf{1}, V_{J_n} \rangle}$$

goes to 0 when n goes to ∞ , with probability 1.

$$\begin{aligned} \epsilon_n &= \frac{\langle f_0, W_n \rangle}{\langle \mathbf{1}, W_n \rangle} - \frac{\langle f_0, V_{J_n} \rangle}{\langle \mathbf{1}, V_{J_n} \rangle} \\ &= \frac{\langle f_0, W_n \rangle - \langle f_0, V_{J_n} \rangle}{\langle \mathbf{1}, W_n \rangle} + \langle f_0, V_{J_n} \rangle \left(\frac{1}{\langle \mathbf{1}, W_n \rangle} - \frac{1}{\langle \mathbf{1}, V_{J_n} \rangle} \right) \\ &= \frac{-\langle f_0, V_{J_n} \ominus W_n \rangle}{\langle \mathbf{1}, W_n \rangle} + \frac{\langle f_0, V_{J_n} \rangle \langle \mathbf{1}, V_{J_n} \ominus W_n \rangle}{\langle \mathbf{1}, V_{J_n} \rangle \langle \mathbf{1}, W_n \rangle}. \end{aligned}$$

We use that $\langle f_0, W_n \rangle \leq \langle \mathbf{1}, W_n \rangle$ for every $n \geq 1$ to get:

$$|\epsilon_n| \leq 2 \frac{\langle \mathbf{1}, V_{J_n} \ominus W_n \rangle}{\langle \mathbf{1}, W_n \rangle} \leq 2 \frac{\Gamma_{\gamma(W_n)}(\omega \ominus W_n)}{\langle \mathbf{1}, W_n \rangle}, \quad (56)$$

where the latter inequality follows from (55). According to Proposition 8.10, and since $(\Gamma_m)_{m \in \mathcal{M}_{rs}}$ is an integrable homogeneous function, the right member in (56) goes to 0 with probability 1, and thus $\lim_{n \rightarrow \infty} \epsilon_n = 0$ with probability 1. This completes the proof of point 1 of Lemma 8.13.

The proof of point 2 follows from a similar calculation. \square

Proof of Theorem 7.8

We shall now prove Theorem 7.8 as a corollary of Theorem 7.6. We also use the result stated in Lemma 8.6, that uses the fact that we consider distributed probabilities.

Let $f = (f(\mathbf{s}, \cdot))_{\mathbf{s} \in \Sigma}$ be any extended state function; we have to study the convergence and the limit of the ergodic means $M(f, V_n)$, for $(V_n)_{n \geq 1}$ any regular sequence of stopping operators. Because of the linearity of $f \rightarrow M(f, \cdot)$, we assume without loss of generality that f satisfies $f(\mathbf{s}, \cdot) = 0$ on $\Omega_{\mathbf{s}}$ whenever $\mathbf{s} \neq \mathbf{s}_0$, where \mathbf{s}_0 is some particular dynamic cluster. Let $(V_n)_{n \geq 0}$ be a regular sequence of stopping operators, and let $\omega \in \Omega$ such that $\lim_{n \rightarrow \infty} M(\mathbf{1}_{\mathbf{s}_0}, V_n(\omega)) = \alpha(\mathbf{s}_0)$ holds. Let also $w_n = S_n^{\mathbf{s}_0}(\omega)$ denote the sequence of local return operators to \mathbf{s}_0 applied to ω . According to Proposition 8.5, we assume without loss of generality that $\sup_{n \geq 1} w_n = \omega$, since the set of such ω s has probability 1.

Recall that the branching cells $x \in \Delta(\omega)$ such that $\langle x \rangle = \mathbf{s}_0$ are totally ordered. Therefore, for every $n \geq 1$, the set of branching cells:

$$I_n = \{x \in \Delta(V_n(\omega)) : \langle x \rangle = \mathbf{s}_0\}$$

is an interval. But since $\sup_{n \geq 1} w_n = \omega$, there is thus an integer $J(n)$ such that:

$$I_n = \{x \in \Delta(w_{J(n)}) : \langle x \rangle = \mathbf{s}_0\}.$$

Since we assume that $f(\mathbf{s}, \cdot)$ vanishes if $\mathbf{s} \neq \mathbf{s}_0$, we have:

$$\begin{aligned} M(f, V_n) &= \frac{1}{\langle \mathbf{1}, V_n \rangle} \langle f, V_n \rangle \\ &= \frac{1}{\langle \mathbf{1}, V_n \rangle} \langle f, w_{J(n)} \rangle \\ &= \frac{\langle \mathbf{1}_{\mathbf{s}_0}, V_n \rangle}{\langle \mathbf{1}, V_n \rangle} \frac{1}{\langle \mathbf{1}_{\mathbf{s}_0}, V_n \rangle} \sum_{k=1}^{J(n)} f(\mathbf{s}_0, Y_k^{\mathbf{s}_0}), \end{aligned} \quad (57)$$

where $Y_k^{\mathbf{s}_0}$ is defined by (39). $\Delta(\omega)$ contains infinitely many instances of \mathbf{s}_0 since the net is recurrent. Since $\sup_n V_n(\omega) = \omega$, it follows that $\lim_{n \rightarrow \infty} J(n) = +\infty$. The sequence $(Y_n^{\mathbf{s}_0})_{n \geq 1}$ is a sequence of independent identically distributed (i.i.d.) random variables of law $p_{\mathbf{s}_0}$ according to Lemma 8.6. It follows thus from the Strong law of large numbers for i.i.d. sequences (see Theorem A.2 in Appendix A) that we have:

$$\begin{aligned} \frac{1}{\langle \mathbf{1}_{\mathbf{s}_0}, V_n \rangle} \sum_{k=1}^{J(n)} f(\mathbf{s}_0, Y_k^{\mathbf{s}_0}) &= \frac{1}{\langle \mathbf{1}_{\mathbf{s}_0}, w_{J(n)} \rangle} \sum_{k=1}^{J(n)} f(\mathbf{s}_0, Y_k^{\mathbf{s}_0}) \\ &= \frac{1}{J(n)} \sum_{k=1}^{J(n)} f(\mathbf{s}_0, Y_k^{\mathbf{s}_0}) \xrightarrow{n \rightarrow \infty} p_{\mathbf{s}_0}(f), \end{aligned} \quad (58)$$

where $p_{\mathbf{s}_0}(f)$ is defined by:

$$p_{\mathbf{s}_0}(f) = \sum_{z \in \Omega_{\mathbf{s}_0}} p_{\mathbf{s}_0}(z) f(\mathbf{s}_0, z).$$

On the other hand, we have

$$\lim_{n \rightarrow \infty} \frac{\langle \mathbf{1}_{s_0}, V_n \rangle}{\langle \mathbf{1}, V_n \rangle} = \alpha(s_0). \quad (59)$$

Combining (57), (58) and (59), we get:

$$\lim_{n \rightarrow \infty} M(f, V_n) = \alpha(s_0)p_{s_0}(f),$$

which is Theorem 7.8 for our particular f . As it was enough to consider this particular f , the proof of the theorem is complete.

9 Related work

This study belongs to the recently developed area of true-concurrency probabilistic models, addressing both event structures and nets. Confusion-free probabilistic event structures were studied in [11, 13]. The viewpoint adopted in [13] is slightly different from ours: they consider probabilistic continuous valuations defined on the domain of configurations of an event structure. So-called non-leaking valuations are of particular interest for them; the later are in bijective correspondence with regular probability measures on the space of maximal configurations, hence both their point of view and ours finally agree. The valuations effectively constructed in [13] are the valuations with independence, for the case of confusion-free event structures. These are built up from the notion of *cell* for confusion-free event structures. The distributed probabilistic event structures we introduced in [9, 6], built up from the more general notion of branching cell, address event structures *with* confusion, and can thus be seen as a generalization of the valuations with independence—although both works were done in parallel. For confusion-free event structures, branching cells reduce to cell, and distributed probabilities are in bijective correspondence with valuations with independence.

The approach of [9, 6] has its origin in [12], where a first model of Markov net were investigated. A study of Markovian properties for concurrent systems is found in [17], while associated topological studies are found in [14, 15].

Probabilistic models have been mostly considered for the purpose of performance evaluation, in relation with queuing theory [19, 20, 21]. For this case, global physical time comes into play, which justifies using classical sequential semantics. In contrast, other applications such as diagnosis mentioned in the Introduction, require a true-concurrency probabilistic setting.

Probabilistic automata, first studied by Rabin [22] are another kind of probabilistic model. When considering products of probabilistic automata, concurrency is taken into account by means of a scheduler [23, 24], which may, or may not, be random itself. When the scheduler is fixed, then probabilistic automata become ordinary Markov chains (with interleaving semantics). The presence of a possibly nondeterministic scheduler make them hybrid models, involving both probabilities and nondeterminism; as such, probabilistic automata belong to the class of Markov Decision Processes. These kind of hybrid models have recently been developed for use in security [25]. We believe that our own work can be of real help in this area.

10 Summary and Perspectives

We have proposed branching cells and dynamic clusters as a notion of local, concurrent, state for event structures and Petri nets. We have applied the construction of distributed probabilities for event structures to the randomization of traces of safe Petri nets. This results in the model of Markov nets, a proper generalization of Markov chains to true-concurrency systems. The Law of Large Numbers extends to Markov nets, with dynamic clusters taken as states.

Although our work relies on the local finiteness assumption for event structures, the notion of branching cells that we have developed can be applied to general event structures arising from Petri net unfoldings. Some new issues arise for non-locally finite event structures, however. First, branching cells may be infinite; second, maximal configurations may not be R -stopped, hence the randomization is not immediate as for locally finite event structures. We are currently working on this extended setting for Markov nets.

A Appendix: Classical Laws of Large Numbers

In this section we state the classical Laws of Large Numbers that we use, namely: for Markov chains and for sequences of independent identically distributed (i.i.d.) random variables.

Theorem A.1 [26, Ch.V, Th. 6.1 p.219] *Let $(X_n)_{n \geq 0}$ be a Markov chain on a finite state space Q , and let $f : Q \rightarrow \mathbb{R}$ be a real-valued function. Then the ergodic means:*

$$\frac{1}{n} \sum_{k=0}^{n-1} f(X_k)$$

have a limit, for $n \rightarrow \infty$, with probability 1.

Theorem A.2 [26, Ch.III, Th. 5.1 p.142] *Let $(X_n)_{n \geq 0}$ be a sequence of i.i.d. random variables, with probability law p on a countable set Q . For any real-valued non-negative function $f : Q \rightarrow \mathbb{R}$, the following limit holds with probability 1:*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) = \sum_{s \in Q} p(s) f(s).$$

Finally, the following result is an exercise:

Proposition A.3 *Let $P = \{p_i, i = 1, \dots, r\}$ be a finite family of probability laws on some countable set Q . Let $(X_n)_{n \geq 0}$ be a sequence of independent random variables, such that, for each $n \geq 0$, the law of X_n belongs to P . Let $f : Q \rightarrow \mathbb{R}$ be a nonnegative real-valued function, such that, for each $p_i \in P$, the following expectation is finite:*

$$q_i(f) = \sum_{q \in Q} p_i(q) f(q) < \infty.$$

Then we have:

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) \leq \max_{i=1, \dots, r} q_i(f) < \infty.$$

ACKNOWLEDGMENT. We gratefully thank an anonymous reviewer for her/his careful reading of the manuscript as well as interesting criticisms and suggestions.

References

- [1] P. Nielsen, G. Plotkin, G. Winskel, Petri nets, event structures and domains, part 1, *Theoretical Computer Science* 13 (1981) 85–108.
- [2] A. Mazurkiewicz, Trace theory, in: W. Brauer, W. Reisig, G. Rozenberg (Eds.), *Petri Nets: Applications and Relationships to Other Models of Concurrency*, Vol. 255 of LNCS, Springer, 1986, pp. 279–324.
- [3] G. Holzmann, D. Peled, An improvement in formal verification, in: D. Hogrefe, S. Leue (Eds.), *FORTE 1994*, Vol. 6 of IFIP Conference Proceedings, Chapman & Hall, 1995, pp. 197–211.
- [4] E. Fabre, A. Benveniste, S. Haar, C. Jard, Distributed monitoring of concurrent and asynchronous systems, *Discrete Event Dynamic Systems* 15 (1) (2005) 33–84.
- [5] E. Fabre, A. Benveniste, S. Haar, C. Jard, Distributed monitoring of concurrent and asynchronous systems, Tech. Rep. RR-4842, INRIA, available at <http://www.inria.fr/rrrt/rr-4842.html> (2004).
- [6] S. Abbes, A. Benveniste, Probabilistic models for true-concurrency: branching cells and distributed probabilities for event structures, *Information and Computation* 204 (2) (2006) 231–274.
- [7] D. Revuz, *Markov Chains*, North Holland, 1975.
- [8] S. Abbes, A. Benveniste, Branching cells as local states for event structures and nets: probabilistic applications, in: V. Sassone (Ed.), *FOSSACS 2005*, Vol. 3441 of LNCS, 2005, pp. 95–109.
- [9] S. Abbes, Probabilistic model for concurrent and distributed systems. limit theorems and application to statistical parametric estimation, Ph.D. thesis, IRISA/Université de Rennes 1 (2004).
- [10] P. Billingsley, *Probability and Measure*, John Wiley, 1986, second edition.
- [11] H. Völzer, Randomized non-sequential processes, in: K. Larsen, M. Nielsen (Eds.), *CONCUR 2001*, Vol. 2154 of LNCS, 2001, pp. 184–201.

- [12] A. Benveniste, S. Haar, E. Fabre, Markov nets: probabilistic models for distributed and concurrent systems, *IEEE Transactions on Automatic Control* 48 (11) (2003) 1936–1950.
- [13] D. Varacca, H. Völzer, G. Winskel, Probabilistic event structures and domains, in: P. Gardner, N. Yoshida (Eds.), *CONCUR 2004*, Vol. 3170 of LNCS, 2004, pp. 481–496.
- [14] S. Abbes, K. Keimel, Projective topology on bifinite domains and applications, *Theoretical Computer Science* 365 (3) (2006) 171–183.
- [15] S. Abbes, A projective formalism for topological and probabilistic event structures, to appear in *Mathematical Structures in Computer Science* (2007).
- [16] S. Abbes, A. Benveniste, Branching cells as local states for event structures and nets: probabilistic applications, Tech. Rep. RR-5347, INRIA, available at <http://www.inria.fr/rrrt/rr-5347.html> (2004).
- [17] S. Abbes, A (true) concurrent Markov property and some applications to Markov nets, in: G. Ciardo, P. Darondeau (Eds.), *ICATPN 2005*, Vol. 3536 of LNCS, 2005, pp. 70–89.
- [18] A. Raugi, J.-P. Conze, Cours de DEA. Martingales et systèmes dynamiques, IRMAR, Université de Rennes 1 (2002).
- [19] L. Kleinrock, *Communication Nets: Stochastic Message Flow and Design*, McGraw-Hill, 1964.
- [20] H. Daduna, *Queuing Networks with Discrete Time Scale*, Vol. 2046 of LNCS, Springer, 2001.
- [21] P. Haas, *Stochastic Petri nets*, Springer Verlag, 2002.
- [22] M. Rabin, Probabilistic automata, *Information and Control* 6 (1963) 230–245.
- [23] R. Segala, A compositional trace-based semantics for probabilistic automata, in: I. Lee, S. Smolka (Eds.), *CONCUR 1995*, Vol. 962 of LNCS, 1995, pp. 234–248.
- [24] N. Lynch, R. Segala, F. Vaandrager, Hybrid I/O automata revisited, in: M. Di Benedetto, A. Sangiovanni-Vincentelli (Eds.), *HSCC 2001*, Vol. 2034 of LNCS, 2001, pp. 403–417.
- [25] S. Umeno, N. Lynch, Proving safety properties of an aircraft landing protocol using I/O Automata and the PVS theorem prover: a case study, in: J. Misra, T. Nipkow, E. Sekerinski (Eds.), *Formal Methods 2006*, Vol. 4085 of LNCS, 2006, pp. 64–80.
- [26] J. Doob, *Stochastic Processes*, John Wiley, 1953.